Correctness and completeness of logic programs

Version of 24-04-2015

Włodzimierz Drabent,
Institute of Computer Science, Polish Academy of Sciences, and IDA, Linköping University, Sweden

arXiv:1412.8739v2 [cs.LO] 17 May 2015

We discuss proving correctness and completeness of definite clause logic programs. We propose a method for proving completeness, while for proving correctness we employ a method which should be well known but is often neglected. Also, we show how to prove completeness and correctness in the presence of SLD-tree pruning, and point out that approximate specifications simplify specifications and proofs.

We compare the proof methods to declarative diagnosis (algorithmic debugging), showing that approximate specifications eliminate a major drawback of the latter. We argue that our proof methods reflect natural declarative thinking about programs, and that they can be used, formally or informally, in everyday programming.

CCS Concepts: •**Theory of computation** → *Program reasoning;* •**Software and its engineering** → *Constraint and logic languages; Software verification; Software testing and debugging;*

General Terms: Theory, Verification

Additional Key Words and Phrases: logic programming, declarative programming, program completeness, program correctness, specifications, declarative diagnosis / algorithmic debugging

1. INTRODUCTION

The purpose of this paper is to present proof methods to deal with the correctness related issues of definite clause logic programs. The goal is to devise methods which are 1. simple and applicable in the practice of Prolog programming, and 2. declarative (i.e. not referring to any operational semantics), in other words depending only on the logical reading of programs.

The notion of program correctness (in imperative and functional programming) divides in logic programming into correctness and completeness. Correctness means that all answers of the program are compatible with the specification, completeness – that the program produces all the answers required by the specification. A specification may be *approximate*: for such a specification some answers are allowed, but not required to be computed. We discuss various advantages of using approximate specifications.

There exists a simple, natural, and moreover complete proof method for correctness [Clark 1979]. It should be well known but has been mostly neglected. Surprisingly little has been done on reasoning about completeness. See Section 8 for the related work. In this paper, for proving correctness we employ Clark's method. For proving completeness we introduce a simplification of the method of [Drabent and Miłkowska 2005]. We introduce a new notion of semi-completeness; semi-completeness and termination imply completeness. We discuss a few sufficient conditions for completeness. We also propose a way of proving that completeness is preserved under pruning of SLD-trees (due to e.g. the cut of Prolog). This is augmented by a way of proving correctness for such trees, taking into account that pruning may remove some answers.

We would like to treat logic programming as a declarative programming paradigm, and to reason about programs declaratively, i.e. independently from their operational semantics. The method of proving correctness is declarative, and so is the method of proving semi-completeness. However some of the sufficient conditions for completeness are not declarative, as they, roughly speaking, refer to an operational notion of program termination. Notice that in most practical cases termination has to be established anyway. So such a way of proving completeness may be a reasonable compromise between declarative and non-declarative reasoning.

This paper provides small examples; for a more substantial application of the presented proof methods see [Drabent 2012a]. That paper presents a construction of a practical Prolog

program, the SAT solver of [Howe and King 2012]. Starting from a formal specification, a logic program is constructed hand in hand with its correctness and completeness proofs. Then control is added (delays, and pruning SLD-trees) to obtain the final Prolog program. The added control preserves correctness, while the approach presented here makes it possible to formally prove that also completeness is preserved. This example illustrates also the usefulness of approximate specifications. The consecutively constructed versions of the program are not equivalent – their main predicates define different relations. So it is not a case of a program transformation which preserves the semantics of a program. What is preserved is correctness and completeness w.r.t. the same approximate specification.

Closely related to proving correctness and completeness of programs is program diagnosis, i.e. locating errors in programs. We show a correspondence between the presented proof methods and declarative diagnosis (called also algorithmic debugging). We show how an important drawback of declarative diagnosis can be easily overcome by employing approximate specifications.

Preliminaries. In this paper we consider definite clause programs (i.e. logic programs without negation). For (standard) notation and definitions the reader is referred to [Apt 1997]. Given a predicate symbol p, by an *atom for* p (or p*-atom*) we mean an atom whose predicate symbol is p, and by a *clause for* p – a clause whose head is an atom for p. The set of the clauses for p in the program under consideration is called *procedure* p.

We assume a fixed alphabet of function and predicate symbols, not restricted to the symbols from the considered program. The only requirement is that the Herbrand universe is nonempty. In particular, the set of function symbols may be finite or infinite. The Herbrand universe will be denoted by $\mathcal{HU}$, the Herbrand base by $\mathcal{HB}$, and the sets of all terms, respectively atoms, by $\mathcal{TU}$ and $\mathcal{TB}$. By $ground(E)$ we mean the set of ground instances of a term, atom, etc E. For a program P, $ground(P)$ is the set of ground instances of the clauses of P. The least Herbrand model of a program P will be denoted by $\mathcal{M}_P$. It is the least fixed point of the immediate consequence operator $T_P\colon 2^{\mathcal{HB}} \to 2^{\mathcal{HB}}$; $\mathcal{M}_P = \bigcup_{i=1}^{\infty} T_P^i(\emptyset)$.

By a computed (respectively correct) answer for a program P and a query Q we mean an instance $Q\theta$ of Q where θ is a computed (correct) answer substitution [Apt 1997] for Q and P. We usually say just *answer* as each computed answer is a correct one, and each correct answer (for Q) is a computed answer (for Q or for some its instance). Thus, by soundness and completeness of SLD-resolution, $Q\theta$ is an answer for P iff $P \models Q\theta$.

By "declarative" (property, reasoning, ...) we mean referring only to logical reading of programs, thus abstracting from any operational semantics. So Q being an answer for P is a declarative property. Properties depending on the order of atoms in clauses will not be considered declarative, as logical reading does not distinguish equivalent formulae, like $\alpha \wedge \beta$ and $\beta \wedge \alpha$.

Names of variables begin with an upper-case letter. We use the list notation of Prolog. So $[t_1, \ldots, t_n]$ $(n \geq 0)$ stands for the list of elements $t_1, \ldots, t_n$. Only a term of this form is considered a list. (Thus terms like $[a, a|X]$, or $[a, a|a]$, where $a \neq [\,]$ is a constant, are not lists). $\mathbb{N}$ stands for the set of natural numbers, and $\mathbb{Z}$ for the set of integers; $f\colon A \hookrightarrow B$ states that f is a partial function from A to B. We use a name "proof tree" instead of "implication tree" of [Apt 1997].

Organization of the paper. The next section introduces the basic notions. Section 3 discusses approximate specifications. Sections 4 and 5 deal with proving, respectively, correctness and completeness of programs. Proving completeness and correctness in the context of pruning of SLD-trees is discussed in Section 6. Section 7 discusses declarative diagnosis. Section 8 is an overview of related work, Section 9 provides additional discussion of the proposed methods, and the next section concludes the paper. Two appendices include results about completeness of the proof methods, and proofs missing in the main text.

2. CORRECTNESS AND COMPLETENESS

This section introduces basic notions of specification, correctness and completeness, and relates them to the answers of programs.

The purpose of a logic program is to compute a relation, or a few relations. It is convenient to assume that the relations are on the Herbrand universe. A specification should describe them. In most cases all the relations defined by a program should be described, an n-ary relation for each n-argument predicate symbol. It is convenient to describe such family of relations as a Herbrand interpretation S: a tuple $(t_1, \ldots, t_n)$ is in the relation corresponding to a predicate symbol p iff $p(t_1, \ldots, t_n) \in S$. So by a (formal) **specification** we mean a Herbrand interpretation; given a specification S, each $A \in S$ is called a *specified atom*.

In imperative programming, correctness usually means that the program results are as specified. In logic programming, due to its non-deterministic nature, we actually have two issues: **correctness** (all the results are compatible with the specification) and **completeness** (all the results required by the specification are produced). In other words, correctness means that the relations defined by the program are subsets of the specified ones, and completeness means inclusion in the opposite direction. In terms of specified atoms and the least Herbrand model $\mathcal{M}_P$ of a program P, we define:

Definition 2.1. Let P be a program and $S \subseteq \mathcal{HB}$ a specification. P is **correct** w.r.t. S when $\mathcal{M}_P \subseteq S$; it is **complete** w.r.t. S when $\mathcal{M}_P \supseteq S$.

We will sometimes skip the specification when it is clear from the context. We propose to call a program **fully correct** when it is both correct and complete.

It is important to relate the answers of programs to their correctness and completeness. The relation is not trivial as, for a program P and a query Q, the equivalence $\mathcal{M}_P \models Q$ iff $P \models Q$ is not true in some cases.[1]

LEMMA 2.2. *Let P be a program, and Q a query such that*

(1) Q contains exactly $k \geq 0$ (distinct) variables, and the underlying language has (at least) k constants not occurring in P, Q, or a non-constant function symbol not occurring in P, Q.

Then $\mathcal{M}_P \models Q$ iff $P \models Q$.

In particular, the equivalence holds for ground queries. It also holds when the set of function symbols is infinite (and P is finite). Note that condition (1) is satisfied by any Prolog program and query, so the equivalence holds in practice. See Appendix B for a proof of the lemma, and [Drabent 2015b] for a slightly more general result and further discussion.

Now we are ready to relate program answers to the correctness and completeness of programs. (Remember that a query Q is an answer for P iff $P \models Q$.)

PROPOSITION 2.3. *Let P be a program, and S a specification.*

(1) P is correct w.r.t. S iff $P \models Q$ implies $S \models Q$ for each query Q.
(2) P is complete w.r.t. S iff $S \models Q$ implies $P \models Q$ for any ground *query Q.*
(3) If the set of function symbols of the underlying language is infinite and P is finite then P is complete w.r.t. S iff $S \models Q$ implies $P \models Q$ for any query Q.
(4) P is complete w.r.t. S iff
$S \models Q$ implies $P \models Q$ for any query Q such that condition (1) of Lemma 2.2 holds.

Within this proposition, "query Q" can be replaced by "atomic query Q".

---

[1] For instance assume a two element alphabet of function symbols, with a unary f and a constant a. Take $P = \{\, p(a).\ p(f(Y)).\}$. Then $\mathcal{M}_P = \mathcal{HU}$. Hence $\mathcal{M}_P \models p(X)$ but $P \not\models p(X)$.

It is quite often believed that the equivalence always holds, See for instance erroneous Th. 2 in [Paschke 2011], or the remark on the least Herbrand models on p. 512 of [Naish 2006].

PROOF. For the first case, assume $\mathcal{M}_P \subseteq S$ (i.e. P is correct). If $P \models Q$ then $\mathcal{M}_P \models Q$, hence $S \models Q$. Conversely, assume that $P \models Q$ implies $S \models Q$ for any atomic query Q. If $A \in \mathcal{M}_P$ then $P \models A$ (by Lemma 2.2), hence $S \models A$. Thus $\mathcal{M}_P \subseteq S$.

For cases 2, 3, 4 it is sufficient to show the "if" of case 2 and the "only-if" of case 4. For the "only-if" of case 4, assume $S \subseteq \mathcal{M}_P$ (P complete) and $S \models Q$ (for a Q as described). Then $\mathcal{M}_P \models Q$ and, by Lemma 2.2, $P \models Q$.

For the "if" of case 2 assume the right hand side of the equivalence, for atomic queries. So for every ground atom A we have that $S \models A$ implies $P \models A$, hence $\mathcal{M}_P \models A$. Thus $S \subseteq \mathcal{M}_P$. □

The following auxiliary notions will be used.

Definition 2.4. The set of p-atoms from a specification S will be called the *specification for* p given by S.

A *predicate* (or a procedure) p in P is *correct* w.r.t. S when each p-atom of $\mathcal{M}_P$ is in S, and *complete* w.r.t. S when each p-atom of S is in $\mathcal{M}_P$.

An *answer* Q is *correct* w.r.t. S when $S \models Q$.

P is *complete for a query* Q w.r.t. S when $S \models Q\theta$ implies that $Q\theta$ is an answer for P, for any ground instance $Q\theta$ of Q.

Informally, complete for Q means that all the answers for Q required by the specification are computed. Note that a program is complete w.r.t. S iff it is complete w.r.t. S for any query iff it is complete w.r.t. S for any query $A \in S$.

There is an alternative way to define program completeness, namely by $S \models Q$ implying $P \models Q$, for every query Q [Drabent and Miłkowska 2005]. It was not chosen here, as the two notions of completeness are equivalent in the practical case of an infinite alphabet of function symbols. Also, it seems that the chosen version leads to simpler sufficient conditions for completeness.

3. APPROXIMATE SPECIFICATIONS

This section first shows that requiring specifications to be exact leads to difficulties and inconveniences. Then usefulness of approximate specifications is discussed.

Notice that if a program P is both correct and complete w.r.t. S then $\mathcal{M}_P = S$ and the specification describes exactly the relations defined by P. Often it is difficult (and not necessary) to specify the relations exactly. The relations defined by the program are often not exactly the ones intended by the programmer. Such discrepancy is however not an error, because for all intended usages of the program the answers are the same as for a program defining the intended relations. For certain atoms it is irrelevant whether they are, or are not, in $\mathcal{M}_P$. We illustrate this by some examples.

Example 3.1. The well-known program

$$\text{APPEND} = \{\ app(\,[H|K], L, [H|M]\,) \leftarrow app(\,K, L, M\,).\quad app(\,[\,], L, L\,).\ \}$$

does not define the relation of list concatenation (provided that the underlying alphabet contains some function symbol not occurring in the program, e.g. a constant a). For instance, APPEND $\models app([\,], a, a)$. In other words, APPEND is not correct w.r.t.

$$S^0_{\text{APPEND}} = \{\, app(k,l,m) \in \mathcal{HB} \mid k, l, m \text{ are lists, } k * l = m \,\},$$

where $k * l$ stands for the concatenation of lists k, l. It is however complete w.r.t. S^0_{APPEND}, and correct w.r.t.

$$S_{\text{APPEND}} = \{\, app(k,l,m) \in \mathcal{HB} \mid \text{if } l \text{ or } m \text{ is a list then } app(k,l,m) \in S^0_{\text{APPEND}} \,\}.$$

Proofs are given later on in Examples 4.3, 5.8. Correctness w.r.t. S_{APPEND} and completeness w.r.t. S^0_{APPEND} are sufficient to show that APPEND will produce the required results when

used to concatenate or split lists.[2] Understanding the exact relation defined by the program is unnecessary. See [Drabent and Miłkowska 2005] for further discussion.

Note that S_{APPEND} is the union of S^0_{APPEND} and the complement $\mathcal{HB} \setminus T$ of the set $T = \{\, app(k,l,m) \in \mathcal{HB} \mid \ l \text{ or } m \text{ is a list} \,\}$. The intuition is that it is irrelevant whether the atoms from $\mathcal{HB} \setminus T$ are in $\mathcal{M}_{\mathrm{APPEND}}$.

When we are only interested in using APPEND to concatenate lists, the specification for correctness can be weakened by replacing "if l or m is a list" by "if l is a list". Similarly, if only list splitting is of interest then "if m is a list" can be used.

A similar example is given in Sect. 7.2 where in an insertion sort program it is irrelevant how elements should be inserted into an unsorted list.

Example 3.2. Programmers often do not know the exact relations defined by logic programs. For many programs in textbooks, the exact relations defined by the predicates are not explained. For instance this concerns the programs in Chapter 3.2 of [Sterling and Shapiro 1994] defining predicates `member/2`, `sublist/2`, `prefix/2`, `suffix/2`, and `append/3`. Similarly to the case of APPEND described above, the relations defined by the programs are not what we understand as the list membership relation, the *sublist* relation etc. The programs define certain supersets of these relations. (They are complete w.r.t. specifications describing the relations.) Suitable specifications for correctness can be easily constructed, similarly as for APPEND. Understanding the exact relations defined by the programs is unnecessary, and they are not discussed in the book.

The exact relations defined by programs are quite often misunderstood. For instance, in [Deville and Lau 1994, Ex. 15] it is claimed that the constructed program, let us call it P, defines the relation of list inclusion

$$\mathcal{I}(included) = \left\{ included(l_1,l_2) \in \mathcal{HB} \;\middle|\; \begin{array}{l} l_1, l_2 \text{ are lists,} \\ \text{every element of } l_1 \text{ belongs to } l_2 \end{array} \right\}.$$

In our terms, this means that predicate *included* of P is correct and complete w.r.t. $\mathcal{I}(included)$. However the correctness does not hold; P contains a clause $included(L_1, L_2) \leftarrow L_1 = [\,]$., thus $P \models included([\,], t)$ for any term t.

The examples show that it is often not necessary to know the semantics of a program exactly, and that its approximate description is sufficient. An **approximate specification** is a pair of specifications S_{compl}, S_{corr} called, respectively, *specification for completeness* and *specification for correctness*. The intention is that the program is complete w.r.t. the former, and correct w.r.t. the latter. We say that a program P is *fully correct* w.r.t. S_{compl}, S_{corr} when $S_{compl} \subseteq \mathcal{M}_P \subseteq S_{corr}$. In other words, the specifications S_{compl}, S_{corr} describe, respectively, which atoms have to be computed, and which are allowed to be computed;

---

[2] Formally: In such usage, in a query $Q = app(s,t,u)$ the argument t (and s) is a list, or u is a list. By the correctness w.r.t. S_{APPEND} and Prop. 2.3(1), for any answer $Q\theta$ of APPEND we have $S_{\mathrm{APPEND}} \models Q\theta$. As $t\theta$ or $u\theta$ is a list, each ground instance of $Q\theta$ is in S^0_{APPEND}. Hence $S^0_{\mathrm{APPEND}} \models Q\theta$. This means that $s\theta, t\theta, u\theta$ are lists and $u\theta$ is the concatenation of $s\theta, t\theta$.

Conversely, let us use the fact that APPEND is complete w.r.t. S^0_{APPEND} for any set of function symbols that includes $[\,]$ and $[\,|\,]$. Whenever $u\sigma$ is a ground list which is the concatenation of $s\sigma$ and $t\sigma$ (equivalently $S^0_{\mathrm{APPEND}} \models Q\sigma$) then, by Prop. 2.3(2), $Q\sigma$ is an answer of APPEND. The same holds for a non-ground list $u\sigma$, provided that the condition of Prop. 2.3(3) or 2.3(4) holds. In particular, when $u = [u_1, \ldots, u_n]$ is a list and s, t are distinct variables then each $app([u_1, \ldots, u_i], [u_{i+1}, \ldots, u_n], u)$ $(i = 0, \ldots, n)$ is an answer of APPEND. In other words, the program produces each splitting of u.

When the conditions of Prop. 2.3 are not satisfied, it is possible that a program P is complete w.r.t. S^0_{APPEND} and $S^0_{\mathrm{APPEND}} \models Q$, but Q is not an answer of P. This happens for instance when $[\,]$ and $[\,|\,]$ are the only function symbols, $Q = app([X], [\,], [X])$, and P consists of three clauses: $app(\,[\,], L, L\,)$.; $app(\,[[\,]|K], L, [[\,]|M]\,) \leftarrow app(\,K, L, M\,)$.; $app(\,[[H|T]|K], L, [[H|T]|M]\,) \leftarrow app(\,K, L, M\,)$. (A completeness proof for P is similar to that of Ex. 5.8).

for the atoms from $S_{corr} \setminus S_{compl}$ the semantics of the program is irrelevant. Sometimes, by abuse of terminology, a single specification S_{corr} or S_{compl} will be called approximate, especially when the intention is that that the specification is distinct from $\mathcal{M}_P$.

Usually programs are intended to be used with queries Q such that if Q is atomic then it has no instances in $S_{corr} \setminus S_{compl}$. For such queries a program fully correct w.r.t. S_{compl}, S_{corr} behaves like a program correct and complete w.r.t. S_{compl}:

PROPOSITION 3.3. *Let S_{compl}, S_{corr} be an approximate specification, P be a program complete w.r.t. S_{compl} and correct w.r.t. S_{corr}, and P' a program correct and complete w.r.t. S_{compl}. Let Q be an atomic query such that $ground(Q) \cap (S_{corr} \setminus S_{compl}) = \emptyset$. Let $Q\theta$ be an instance of Q such that condition (1) of Lemma 2.2 holds for P with $Q\theta$ and for P' with $Q\theta$. Then $Q\theta$ is an answer for P iff $Q\theta$ is an answer for P'.*

In particular, the equivalence holds when the set of function symbols is infinite (and P, P' are finite).

PROOF. For any program P'' such that P'' and Q satisfy (1), it holds that $P'' \models Q\theta$ iff $ground(Q\theta) \subseteq M_{P''}$. Note that $\mathcal{M}_{P'} = S_{compl} \subseteq \mathcal{M}_P \subseteq S_{corr}$. We have $ground(Q) \cap S_{corr} \subseteq S_{compl}$. Hence $ground(Q\theta) \cap \mathcal{M}_P \subseteq S_{compl}$. Thus $ground(Q\theta) \subseteq \mathcal{M}_P$ iff $ground(Q\theta) \subseteq S_{compl}$. Hence $P \models Q\theta$ iff $P' \models Q\theta$. □

Comments. An obvious role of approximate specifications is to describe selected properties of programs. Ex. 4.2 below shows a case where only the lengths of lists are taken into account, not their elements. Note that such a property cannot be expressed by an exact specification.

Typed logic programming [Pfenning 1992] is outside of the scope of this paper. But it is worth mentioning here that, for some programs, introducing types may change their semantics, so that they define exactly the intended relations. For instance in some appropriate typed logic, the APPEND program of Ex. 3.1 (with the arguments of *app* typed to be lists) would define exactly the relation of list concatenation Thus such typed APPEND would be correct and complete w.r.t. S^0_{APPEND}, and the need for an approximate specification disappears. On the other hand, to obtain a similar result for procedure *insert* of Ex. 7.2 it is necessary to employ a type of sorted integer lists, which seems to be outside of the scope of known type systems for logic programs.

Often, while developing a program, the exact relations to be defined should not be fixed in advance. This would mean taking some design decisions too early. For instance, constructing an append program with S^0_{APPEND} (or with S_{APPEND}) as an exact specification would lead to a program less efficient and more complicated than APPEND. See also Ex. 7.2. A larger example is shown in [Drabent 2012a; 2012b]: a program is derived by stepwise refinement, and the semantics of (the common predicates in) the consecutive versions of the program differ. What does not change is an approximate specification (of the common predicates); the correctness and completeness is preserved in all the versions.

This example suggests a generalization of the paradigm of program development by transformations which preserve program semantics [Pettorossi et al. 2010, and the references therein]. It shows that it is useful and natural to use transformations which instead preserve correctness and completeness w.r.t. an approximate specification.

Importance of approximate specifications for declarative diagnosis is discussed in Sect. 7.

4. PROVING CORRECTNESS

To prove correctness we use the following property [Clark 1979].

THEOREM 4.1 (CORRECTNESS). *A sufficient condition for a program P to be correct w.r.t. a specification S is*

$$S \models P.$$

So the sufficient condition states that, for each ground instance $H \leftarrow B_1, \ldots, B_n$ $(m \geq 0)$ of a clause of the program, if $B_1, \ldots, B_n \in S$ then $H \in S$. Equivalently, $T_P(S) \subseteq S$.

PROOF. $S \models P$ means that S is a Herbrand model of P, thus $\mathcal{M}_P \subseteq S$. □

Example 4.2. Consider a program SPLIT:

$$s([\,],[\,],[\,]). \tag{1}$$
$$s([X|Xs],[X|Ys],Zs) \leftarrow s(Xs,Zs,Ys). \tag{2}$$

Let $|l|$ denote the length of a list l. A specification

$$S = \{\, s(l,l_1,l_2) \mid l,l_1,l_2 \text{ are lists, } 0 \leq |l_1| - |l_2| \leq 1 \,\}$$

describes how the sizes of the last two arguments of S are related. To show, by Th. 4.1, that the program is correct w.r.t. S, consider a ground instance $s([h|t],[h|t_2],t_1) \leftarrow s(t,t_1,t_2)$ of (2). Assume $s(t,t_1,t_2) \in S$. Thus $[h|t],[h|t_2],t_1$ are lists. Let $m = |t_1| - |t_2|$. As $m \in \{0,1\}$, we have $|[h|t_2]| - |t_1| = 1 - m \in \{0,1\}$. So the head $s([h|t],[h|t_2],t_1)$ is in S. The proof for (1) is trivial.

A stronger specification with respect to which SPLIT is correct is shown below.

Example 4.3. To show that program APPEND is correct w.r.t. S_{APPEND} from Ex. 3.1, take a $(H \leftarrow B) \in \mathit{ground}(\text{APPEND})$ such that $B = \mathit{app}(k,l,m) \in S_{\text{APPEND}}$. So $H = \mathit{app}([h|k],l,[h|m])$. If l or $[h|m]$ is a list then k,l,m are lists and $k * l = m$. So $[h|k],l,[h|m]$ are lists and $[h|k] * l = [h|m]$. Thus $H \in S_{\text{APPEND}}$.

For the other clause of the program, take $\mathit{app}([\,],m,m) \in \mathit{ground}(\text{APPEND})$. If m is a list then $[\,] * m = m$; so $\mathit{app}([\,],m,m) \in S_{\text{APPEND}}$).

Example 4.4. Now we show correctness of SPLIT w.r.t. a more precise specification

$$S_{\text{SPLIT}} = \left\{ \begin{array}{l} s([t_1,\ldots,t_{2n}],[t_1,\cdots,t_{2n-1}],[t_2,\cdots,t_{2n}]), \\ s([t_1,\ldots,t_{2n+1}],[t_1,\cdots,t_{2n+1}],[t_2,\cdots,t_{2n}]) \end{array} \;\middle|\; \begin{array}{l} n \geq 0, \\ t_1,\ldots,t_{2n+1} \in \mathcal{HU} \end{array} \right\},$$

where $[t_k,\cdots,t_l]$ denotes the list $[t_k,t_{k+2},\ldots,t_l]$, for k,l being both odd or both even. (When $k > l$, $[t_k,\cdots,t_l] = [t_k,\ldots,t_l] = [\,]$.) For (1) the proof is trivial. Consider a ground instance $H \leftarrow B$ of (2). Assume $B \in S_{\text{SPLIT}}$. If $B = s([t_1,\ldots,t_{2n}],[t_1,\cdots,t_{2n-1}],[t_2,\cdots,t_{2n}])$ then $H = s([t_0,t_1,\ldots,t_{2n}],[t_0,t_2,\cdots,t_{2n}],[t_1,\cdots,t_{2n-1}])$, for some t_0. If $B = s([t_1,\ldots,t_{2n+1}],[t_1,\cdots,t_{2n+1}],[t_2,\cdots,t_{2n}])$ then $H = s([t_0,t_1,\ldots,t_{2n+1}],[t_0,t_2,\cdots,t_{2n}],[t_1\cdots,t_{2n+1}])$. In both cases $H \in S_{\text{SPLIT}}$. (To see this, rename each t_i as t'_{i+1}.)

Completeness of this correctness proving method is discussed in Appendix A. For more examples (among others involving difference lists and accumulators), and for further explanations, references and discussion see [Drabent and Miłkowska 2005].

5. PROVING COMPLETENESS

We first introduce a notion of semi-completeness, and sufficient conditions under which semi-completeness of a program implies its completeness. The next subsection presents a sufficient condition for semi-completeness. For cases where completeness does not follow from semi-completeness, we present a way of proving completeness directly.

5.1. Semi-completeness

Before discussing semi-completeness we present a few auxiliary notions, which have been introduced in a context of proving program termination.

Definition 5.1. A **level mapping** is a function $|\ |\colon \mathcal{HB} \to \mathbb{N}$ assigning natural numbers to atoms.

A program P is **recurrent** *w.r.t. a level mapping* $|\ |$ [Bezem 1993; Apt 1997] if, in every ground instance $H \leftarrow B_1, \ldots, B_n \in \mathit{ground}(P)$ of its clause ($n \geq 0$), $|H| > |B_i|$ for all $i = 1, \ldots, n$. A program is *recurrent* if it is recurrent w.r.t. some level mapping.

Definition 5.2. A program P is **acceptable** w.r.t. a specification S and a level mapping $|\ |$ if P is correct w.r.t. S, and for every $H \leftarrow B_1, \ldots, B_n \in \mathit{ground}(P)$ we have $|H| > |B_i|$ whenever $S \models B_1, \ldots, B_{i-1}$. A program is *acceptable* if it is acceptable w.r.t. some level mapping and some specification.

This definition is more general than the original one [Apt and Pedreschi 1993; Apt 1997] which uses a model of P, instead of a specification for which P is correct. Both definitions are equivalent.[3]

We now introduce a notion of semi-completeness. It may be seen as a step on the way to proving completeness.

Definition 5.3. A program P is **semi-complete** w.r.t. a specification S if P is complete w.r.t. S for (cf. Df. 2.4) any query Q for which there exists a finite SLD-tree.

In practice, the existence of a finite SLD-tree means that P with Q terminates under some selection rule. For a semi-complete program, if a computation for a query Q terminates then it has produced all the answers for Q required by the specification. Obviously, a complete program is semi-complete. We also have:

Proposition 5.4 (Completeness). *Let a program P be semi-complete w.r.t. S. The program is complete w.r.t S if*

(1) for each query $A \in S$ there exists a finite SLD-tree, or
(2) the program is recurrent, or
(3) the program is acceptable (w.r.t. a specification S' possibly distinct from S).

Note that condition (1) is equivalent to each $A \in S$ being an instance of a query Q with a finite SLD-tree.

Proof. For a program P semi-complete w.r.t. S, the first part of condition (1) implies that P is complete w.r.t. S for each query $A \in S$; hence $S \subseteq \mathcal{M}_P$. Condition (2), or (3) implies that for any ground query each SLD-tree, respectively each LD-tree, is finite [Apt 1997]. This implies (1). □

5.2. Proving semi-completeness

We need the following notion.

Definition 5.5. A ground atom H is **covered** *by a clause* C w.r.t. a specification S if H is the head of a ground instance $H \leftarrow B_1, \ldots, B_n$ ($n \geq 0$) of C, such that all the atoms $B_1, \ldots, B_n$ are in S.

A ground atom H is *covered by a program* P w.r.t. S if it is covered w.r.t. S by some clause $C \in P$. A specification S is *covered by a program* P if each $H \in S$ is covered by P w.r.t. S.

For instance, given a specification $S = \{\, p(s^i(0)) \mid i \geq 0 \,\}$, atom $p(s(0))$ is covered both by $p(s(X)) \leftarrow p(X)$ and by $p(X) \leftarrow p(s(X))$. The notion of an atom covered by a clause stems from [Shapiro 1983]. A specification S covered by P is called in [Apt 1993] a supported Herbrand interpretation of P.

---

[3] They define the same programs to be acceptable. More precisely, a program P is acceptable w.r.t. $|\ |$ and a specification S iff, according to the original definition, P is acceptable w.r.t. $|\ |$ and a model I of P. (The "only-if" case holds for each model I such that $\mathcal{M}_P \subseteq I \subseteq S$, the "if" case for each S such that $\mathcal{M}_P \subseteq S \subseteq I$.)

Now we are ready to present a sufficient condition for semi-completeness, which together with Prop. 5.4 provides a sufficient condition for completeness.

THEOREM 5.6 (SEMI-COMPLETENESS). *If a specification S is covered by a program P then P is semi-complete w.r.t. S.*

The proof is presented in Appendix B. Note that the sufficient condition of Th. 5.6 is equivalent to $S \subseteq T_P(S)$; the latter implies $S \subseteq \mathit{gfp}(T_P)$. It is also equivalent to S being a model of ONLY-IF(P), a theory used in defining the completion of P (see e.g. [Doets 1994] or [Drabent and Miłkowska 2005] for a definition).

Example 5.7. We show that program SPLIT from Ex. 4.2 is complete w.r.t. the specification from Ex. 4.4,

$$S_{\mathrm{SPLIT}} = \left\{ \begin{array}{l} s([t_1,\ldots,t_{2n}],[t_1,\cdots,t_{2n-1}],[t_2,\cdots,t_{2n}]), \\ s([t_1,\ldots,t_{2n+1}],[t_1,\cdots,t_{2n+1}],[t_2,\cdots,t_{2n}]) \end{array} \;\middle|\; \begin{array}{l} n \geq 0, \\ t_1,\ldots,t_{2n+1} \in \mathcal{HU} \end{array} \right\},$$

where $[t_k,\cdots,t_l]$ denotes the list $[t_k,t_{k+2},\ldots,t_l]$, for k,l being both odd or both even.

Atom $s([\,],[\,],[\,]) \in S_{\mathrm{SPLIT}}$ is covered by clause (1). For $n > 0$, any atom $A = s([t_1,\ldots,t_{2n}],[t_1,\cdots,t_{2n-1}],[t_2,\cdots,t_{2n}])$ is covered by an instance of (2) with a body $B = s([t_2,\ldots,t_{2n}],[t_2,\cdots,t_{2n}],[t_3,\cdots,t_{2n-1}])$. Similarly, for $n \geq 0$ and any atom $A = s([t_1,\ldots,t_{2n+1}],[t_1,\cdots,t_{2n+1}],[t_2,\cdots,t_{2n}])$, the corresponding body is $B = s([t_2,\ldots,t_{2n+1}],[t_2,\cdots,t_{2n}],[t_3,\cdots,t_{2n+1}])$. In both cases, $B \in S_{\mathrm{SPLIT}}$. (To see this, rename each t_i as t'_{i-1}.) So S_{SPLIT} is covered by SPLIT. Thus SPLIT is semi-complete w.r.t. S_{SPLIT}, by Th. 5.6.

Now by Prop. 5.4 the program is complete, as it is recurrent under the level mapping

$$|\,s(t,t_1,t_2)| = |t|, \quad \text{where} \quad \begin{array}{l} |\,[h|t]\,| = 1 + |t|, \\ |f(t_1,\ldots,t_n)| = 0 \;\; \text{where } n \geq 0 \text{ and } f \text{ is not } [\;|\;], \end{array}$$

(for any ground terms $h,t,t_1,\ldots,t_n$, and any function symbol f).

Example 5.8. Consider program APPEND and specification S^0_{APPEND} from Ex. 3.1. We encourage the reader to check that each $A \in S^0_{\mathrm{APPEND}}$ is covered (by APPEND w.r.t. S^0_{APPEND}). Hence by Th. 5.6 the program is semi-complete (w.r.t. S^0_{APPEND}). Now its completeness follows by Prop. 5.4, as the program is recurrent under the level mapping $|\mathit{app}(t,t_1,t_2)| = |t|$, where $|t|$ is as in the Ex. 5.7.

The notion of semi-completeness is tailored for finite programs. An SLD-tree for an infinite program may be infinite, but with all branches finite. Then a slight generalization of Th. 5.6 applies (Appendix, Prop. B.2, Ex. B.3).

5.3. Proving completeness directly

Semi-completeness and completeness are declarative properties of programs. However the sufficient conditions for completeness of Prop. 5.4(1) and 5.4(3) are not declarative. The former refers explicitly to an operational notion of SLD-tree. The latter employs the notion of acceptable program, which depends on the order of atoms in program clauses, and is related to LD-resolution; this should not be considered as a purely logical reading of formulae. Only condition (2) of Prop. 5.4 can be considered declarative, as it solely depends on logical reading of program clauses. However, this condition is often inapplicable, as many practical programs are not recurrent. Note that the sufficient condition for semi-completeness of Th. 5.6 is declarative.

We are interested in declarative reasoning about logic programs. So we present a declarative sufficient condition for completeness. Moreover, it is applicable in cases where completeness does not follow from semi-completeness. We first generalize the notion of level

mapping by allowing that it is a *partial* function $|\ |\colon \mathcal{HB} \hookrightarrow \mathbb{N}$ assigning natural numbers to some atoms.

Definition 5.9. A ground atom H is **recurrently covered** by a program P w.r.t. a specification S and a level mapping $|\ |\colon \mathcal{HB} \hookrightarrow \mathbb{N}$ if H is the head of a ground instance $H \leftarrow B_1, \ldots, B_n$ $(n \geq 0)$ of a clause of the program, such that $|H|, |B_1|, \ldots |B_n|$ are defined, $B_1, \ldots, B_n \in S$, and $|H| > |B_i|$ for all $i = 1, \ldots, n$.

For instance, given a specification $S = \{\, p(s^i(0)) \mid i \geq 0 \,\}$, atom $p(s(0))$ is recurrently covered by a program $\{\, p(s(X)) \leftarrow p(X).\ \ p(X) \leftarrow p(X).\}$ under a level mapping for which $|p(s^i(0))| = i$. No atom is recurrently covered by a program $\{\, p(X) \leftarrow p(X).\}$. Obviously, if H is recurrently covered by P then it is covered by P.

The following theorem is a reformulation of Th. 6.1 from [Deransart and Małuszyński 1993] (derived there within a sophisticated theory relating attribute grammars and logic programs).

Theorem 5.10 (Completeness 2). *If, under some level mapping* $|\ |\colon \mathcal{HB} \hookrightarrow \mathbb{N}$*, all the atoms from a specification* S *are recurrently covered by a program* P *w.r.t.* S *then* P *is complete w.r.t.* S*.*

Proof. Note first that $|A|$ is defined for all $A \in S$. We show that $A \in \mathcal{M}_P$ for each $A \in S$, by induction on $|A|$. Let $n \geq 0$, assume that $B \in \mathcal{M}_P$ whenever $B \in S$ and $|B| < n$. Take an $A \in S$ with $|A| = n$. As A is recurrently covered, there is a clause instance $A \leftarrow B_1, \ldots, B_m$ $(m \geq 0)$ with $B_i \in S$ and $|B_i| < n$ for each i. So $B_1, \ldots, B_m \in \mathcal{M}_P$, and $A \in \mathcal{M}_P$. □

Example 5.11. Consider a directed graph E. As a specification for a program describing reachability in E, take $S = S_p \cup S_e$, where $S_p = \{\, p(t,u) \mid$ there is a path from t to u in $E \,\}$, and $S_e = \{\, e(t,u) \mid (t,u)$ is an edge in $E \,\}$. Let P consist of a procedure p:

$$p(X,X).$$
$$p(X,Z) \leftarrow e(X,Y),\ p(Y,Z).$$

and a procedure e, which is a set of unary clauses describing the edges of the graph. Assume the latter is complete w.r.t. S_e. Notice that when E has cycles then infinite SLD-trees cannot be avoided, and completeness of P cannot be shown by Prop. 5.4.

To apply Th. 5.10, let us define a level mapping for the elements of S such that $|e(t,u)| = 0$ and $|p(t,u)|$ is the length of a shortest path in E from t to u (so $|p(t,t)| = 0$). Consider an atom $p(t,u) \in S$ where $t \neq u$. Let $t = t_0, t_1, \ldots, t_n = u$ be a shortest path from t to u. Then $e(t,t_1), p(t_1,u) \in S$, $|p(t,u)| = n$, $|e(t,t_1)| = 0$, and $|p(t_1,u)| = n-1$. Thus $p(t,u)$ is recurrently covered by P w.r.t. S and $|\ |$. The same trivially holds for the remaining atoms of S, as they are instances of unary clauses of P. By Th. 5.10, P is complete w.r.t. S.

Note that P is not correct w.r.t. S, unless $\mathcal{HU}$ is the set of nodes of E (as $p(t,t) \notin S$ when t is not a node of the graph). This, again, illustrates usefulness of approximate specifications.

Completeness of the presented completeness proof methods is discussed in Appendix A. In particular it is shown that, in a certain sense, the sufficient condition for semi-completeness of Th. 5.6 is a necessary condition for completeness.

6. PRUNING

Pruning some parts of SLD-trees is often used to improve efficiency of programs. In Prolog it is implemented by using the cut, the if-then-else construct, or by built-ins, like `once/1`. Pruning preserves the correctness of a logic program, it also preserves termination under a given selection rule, but may violate the program's completeness. Consequently, the answers of a program P executed with pruning may be correct w.r.t. a specification S for which P is

not correct; a related programming technique is known as "red cuts" [Sterling and Shapiro 1994].

This section deals with reasoning about program answers in the presence of pruning. The focus is on proving completeness. We begin with a formalization of pruning. Then we introduce a sufficient condition for completeness, and present examples. The next subsection identifies a class of cases where the sufficient condition does not hold (despite pruned trees being complete), and shows how a completeness proof can be obtained anyway. Then we present a sufficient condition for completeness dealing also with infinite pruned trees. We conclude with an approach to prove correctness of the answers of pruned trees. In some examples it is assumed that the reader is familiar with the operational semantics of the cut in Prolog.

6.1. Pruned SLD-trees

By a **pruned SLD-tree** for a program P and a query Q we mean a tree with the root Q which is a connected subgraph of an SLD-tree for P and Q. By an *answer* of a pruned SLD-tree we mean the computed answer of a successful SLD-derivation which is a branch of the tree. We will say that a pruned SLD-tree T with root Q is **complete** w.r.t. a specification S if, for any ground $Q\theta$, $S \models Q\theta$ implies that $Q\theta$ is an instance of an answer of T. Informally, such a tree produces all the answers for Q required by S.

To facilitate reasoning about the answers of pruned SLD-trees, we will view pruning as applying only certain clauses while constructing the children of a node. So we introduce subsets $\Pi_1, \ldots, \Pi_n$ of P. The intention is that for each node the clauses of one Π_i are used. Programs $\Pi_1, \ldots, \Pi_n$ may be not disjoint.

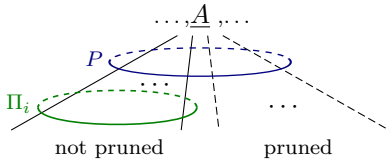

Definition 6.1. Given programs $\Pi_1, \ldots, \Pi_n$ $(n > 0)$, a **c-selection rule** is a function assigning to a query Q' an atom A in Q' and one of the programs $\emptyset, \Pi_1, \ldots, \Pi_n$.

A **csSLD-tree** (cs for clause selection) for a query Q and programs $\Pi_1, \ldots, \Pi_n$, via a c-selection rule R, is constructed as an SLD-tree, but for each node its children are constructed using the program selected by the c-selection rule. An *answer* of the csSLD-tree is the answer of a successful derivation which is a branch of the tree.

A c-selection rule may choose the empty program, thus making a given node a leaf. Note that program Π_i selected for a node Q' may contain clauses whose heads do not unify with the selected atom of Q' (thus inapplicable in constructing the children of the node). A csSLD-tree for Q and $\Pi_1, \ldots, \Pi_n$ is a pruned SLD-tree for Q and $\bigcup_i \Pi_i$. Conversely, for each pruned SLD-tree T for Q and a (finite) program P there exist $n > 0$, and $\Pi_1, \ldots, \Pi_n \subseteq P$ such that T is a csSLD-tree for Q and $\Pi_1, \ldots, \Pi_n$.

It may be needed that different atoms or programs are selected for repeated occurrences of the same query. So, formally, a c-selection rule is a function of a sequence of queries (a path in a tree) $Q, \ldots, Q'$. Also, strictly speaking, not an atom, but an occurrence of an atom in Q' is selected. In some cases it may be convenient to treat a c-selection rule as a function of not only the sequence of queries, but also of (the not pruned part of) the subtree rooted at Q' (as often the decision about pruning is taken after a part of the subtree has been computed).

Example 6.2. Consider a Prolog program

```
c(L1,L2) :- m(X,L1),m(X,L2),!.        m(X,[X|_]).        m(X,[_|L]) :- m(X,L).
```

which is a logic program P with the cut added. Let $\Pi_1 = P$ and $\Pi_2 = \{m(X,[X|_]).\}$. Consider a query $Q = c([t_1,\ldots,t_k],[u_1,\ldots,u_n])$, the SLD-tree T for Q via the Prolog selection rule, and the csSLD-tree T' which is T pruned due to the cut. If Q does not succeed then $T' = T$. Otherwise T' is T without all the nodes to the right of the leftmost successful branch. So in T', for a node Q' with an m-atom selected, Π_2 is selected whenever there is a success in the subtree rooted in the first child of Q'; otherwise Π_1 is selected. Note that the single cut may prune children of many nodes of T.

The c-selection rule of T' can be expressed as a function of a node: Π_2 is selected for a node $Q_i = m(X,[t_i,\ldots,t_k]), m(X,[u_1,\ldots,u_n])$ in T' if t_i is unifiable with some element of $[u_1,\ldots,u_n]$. Otherwise Π_1 is selected for Q_i. For a node of the form $Q'_j = m(t,[u_j,\ldots,u_n])$, program Π_2 is selected if t is unifiable with u_j, otherwise Π_1 is selected.

In this example, determining the c-selection rule corresponding to the cut was not straightforward. Simpler cases are provided by Examples 6.6, 6.7, 6.8 below.

6.2. Completeness of pruned trees

We begin with an example illustrating some of the difficulties related to completeness of pruned csSLD-trees. Then, after introducing some auxiliary notions, a sufficient condition for completeness is presented. A few example proofs follow. We conclude with an example of a completeness proof in a case for which the sufficient condition is not applicable directly.

Example 6.3. We show that completeness of each of $\Pi_1,\ldots,\Pi_n$ is not sufficient for completeness of a csSLD-tree for $\Pi_1,\ldots,\Pi_n$. Consider a program P:

$$q(X) \leftarrow p(Y,X). \quad (3) \qquad p(a,s(X)) \leftarrow p(a,X). \quad (5)$$
$$p(Y,0). \quad (4) \qquad p(b,s(X)) \leftarrow p(b,X). \quad (6)$$

and programs $\Pi_1 = \{(3),(4),(6)\}$, $\Pi_2 = \{(3),(4),(5)\}$, As a specification for completeness consider $S_0 = \{\, q(s^j(0)) \mid j \geq 0 \,\}$. Each of the programs Π_1, Π_2, P is complete w.r.t. S_0. Assume a c-selection rule R choosing alternatively Π_1, Π_2 along each branch of a tree. Then the csSLD-tree for $q(s^j(0)) \in S_0$ via R (where $j > 2$) has no answers, thus the tree is not complete w.r.t. S_0. Similarly the csSLD-tree for $q(X)$ is not complete.

Consider programs $P, \Pi_1,\ldots,\Pi_n$ and specifications $S, S_1,\ldots,S_n$, such that $P \supseteq \bigcup_{i=1}^n \Pi_i$ and $S = \bigcup_{i=1}^n S_i$. The intention is that each S_i describes which answers are to be produced by using Π_i in the first resolution step. We will call $\Pi_1,\ldots,\Pi_n$, $S_1,\ldots,S_n$ a **split** (of P and S).

Definition 6.4. Let $\mathcal{S} = \Pi_1,\ldots,\Pi_n$, $S_1,\ldots,S_n$ be a split, and $S = \bigcup S_i$. Specification S_i is **suitable** for an atom A w.r.t. $\mathcal{S}$ when no instance of A is in $S \setminus S_i$. (In other words, when $ground(A) \cap S \subseteq S_i$.) We also say that a program Π_i is **suitable** for A w.r.t. $\mathcal{S}$ when S_i is.

A c-selection rule is **compatible** with $\mathcal{S}$ if for each non-empty query Q' it selects an atom A and a program Π, such that

– $\Pi \in \{\Pi_1,\ldots,\Pi_n\}$ is suitable for A w.r.t. $\mathcal{S}$, or
– none of $\Pi_1,\ldots,\Pi_n$ is suitable for A w.r.t. $\mathcal{S}$ and $\Pi = \emptyset$ (so Q' is a leaf).

A csSLD-tree for $\Pi_1,\ldots,\Pi_n$ via a c-selection rule compatible with $\mathcal{S}$ is said to be **weakly compatible** with $\mathcal{S}$. The tree is **compatible** with $\mathcal{S}$ when for each its nonempty node some Π_i is selected.

The intuition is that when Π_i is suitable for A then S_i is a fragment of S sufficient to deal with A. It describes all the answers for query A required by S.

The reason of incompleteness of the trees in Ex. 6.3 may be understood as selecting a Π_i not suitable for the selected atom. Take $\mathcal{S} = \Pi_1, \Pi_2, S_0 \cup S_1', S_0 \cup S_2'$, where $S_1' = \{\, p(b, s^i(0)) \mid i \geq 0 \,\}$ and $S_2' = \{\, p(a, s^i(0)) \mid i \geq 0 \,\}$. In the incomplete trees, Π_1 is selected for an atom $A = p(a, u)$, or Π_2 is selected for an atom $B = p(b, u)$ (where $u \in \mathcal{TU}$). However Π_1 is not suitable for A whenever A has an instance in S (as then $\mathit{ground}(A) \cap S \subseteq S_2'$, thus $\mathit{ground}(A) \cap S \not\subseteq S_0 \cup S_1'$); similarly for Π_2 and B.

When Π_i is suitable for A then if each atom of S_i is covered by Π_i (w.r.t. S) then using for A only the clauses of Π_i does not impair completeness w.r.t. S:

THEOREM 6.5. *Let $P \supseteq \bigcup_{i=1}^{n} \Pi_i$ (where $n > 0$) be a program, $S = \bigcup_{i=1}^{n} S_i$ a specification, and T a csSLD-tree for $\Pi_1, \ldots, \Pi_n$. If*

(1) for each $i = 1, \ldots, n$, all the atoms from S_i are covered by Π_i w.r.t. S, and
(2) T is compatible with $\Pi_1, \ldots, \Pi_n, S_1, \ldots, S_n$,
(3) (a) T is finite, or
(b) P is recurrent, or
(c) P is acceptable and T is built under the Prolog selection rule

then T is complete w.r.t. S.

See Appendix B for a proof. Note that in 3c the program may be acceptable w.r.t. a specification distinct from S.

Example 6.6. A Prolog program [Clocksin and Mellish 1981]

```
nop(adam,0) :- !.    nop(eve,0) :- !.    nop(X,2).
```

is an example of difficulties and dangers of using the cut in Prolog. (For its (in)correctness, see Ex. 6.14.) Due to the two occurrences of the cut, for an atomic query A only the first clause with the head unifiable with A will be used. The program can be seen as logic program $P = \Pi_1 \cup \Pi_2 \cup \Pi_3$ executed with pruning, where (for $i = 1, 2, 3$) Π_i is the i-th clause of the program with the cut removed. The intended meaning is $S = S_1 \cup S_2 \cup S_3$, where $S_1 = \{nop(adam, 0)\}$, $S_2 = \{nop(eve, 0)\}$, and $S_3 = \big\{nop(t, 2) \in \mathcal{HB} \mid t \notin \{adam, eve\}\big\}$. Note that all the atoms from S_i are covered by Π_i (for $i = 1, 2, 3$).

Let $\mathcal{S}$ be $\Pi_1, \Pi_2, \Pi_3, S_1, S_2, S_3$. Consider a query $A = nop(t, Y)$ with a ground t. If $t = adam$ then only Π_1 is suitable for A w.r.t. $\mathcal{S}$, as $\mathit{ground}(A) \cap S = S_1$. If $t = eve$ then only Π_2 is suitable. For $t \notin \{adam, eve\}$ the suitable program is Π_3. So for a query A the pruning due to the cuts in the program results in selecting a suitable Π_i, and the obtained csSLD-tree is compatible with $\mathcal{S}$. By Th. 6.5 the tree is complete w.r.t. S.

For a query $nop(X, Y)$ or $nop(X, 0)$ only the first clause, i.e. Π_1, is used. However Π_1 is not suitable for the query (w.r.t. $\mathcal{S}$), and the csSLD-tree is not compatible with $\mathcal{S}$. The premises of Th. 6.5 do not hold, and the tree is not complete (w.r.t. S).

Example 6.7. The following program SAT0 is a simplification of a fragment of the SAT solver of [Howe and King 2012] discussed in [Drabent 2012a; 2012b]. Pruning is crucial for the efficiency and usability of the original program.

$$p(P\text{-}P, [\,]). \quad (7) \qquad p(V\text{-}P, [B|T]) \leftarrow q(V\text{-}P, [B|T]). \quad (8) \qquad p(V\text{-}P, [B|T]) \leftarrow q(B, [V\text{-}P|T]). \quad (9)$$

$$q(V\text{-}P, _) \leftarrow V = P. \quad (10) \qquad q(_, [A|T]) \leftarrow p(A, T). \quad (11) \qquad P = P. \quad (12)$$

The program is complete w.r.t. a specification

$$S = \left\{ \begin{array}{l} p(t_0\text{-}u_0, [t_1\text{-}u_1, \ldots, t_n\text{-}u_n]), \\ q(t_0\text{-}u_0, [t_1\text{-}u_1, \ldots, t_n\text{-}u_n]) \end{array} \;\middle|\; \begin{array}{l} n \geq 0,\ t_0, \ldots, t_n, u_0, \ldots, u_n \in \mathbb{T}, \\ t_i = u_i \text{ for some } i \in \{0, \ldots, n\} \end{array} \right\} \cup S_=$$

where $\mathbb{T} = \{false, true\} \subseteq \mathcal{HU}$, and $S_= = \{\, t{=}t \mid t \in \mathcal{HU} \,\}$. We omit a completeness proof, mentioning only that SAT0 is recurrent w.r.t. a level mapping $|p(t,u)| = 2|u| + 2$, $|q(t,u)| = 2|u| + 1$, $|{=}(t,u)| = 0$, where $|u|$ for $u \in \mathcal{HU}$ is defined in Ex. 5.7.

The first case of pruning is related to redundancy within (8), (9); both $\Pi_1 = \text{SAT0} \setminus \{(9)\}$ and $\Pi_2 = \text{SAT0} \setminus \{(8)\}$ are complete w.r.t. S. For any selected atom at most one of (8), (9) is used, and the choice is dynamic. As the following reasoning is independent from this choice, we omit further details.

So in such pruned SLD-trees the children of each node are constructed using one of programs Π_1, Π_2. Thus they are csSLD-trees for Π_1, Π_2. They are compatible with $\mathcal{S} = \Pi_1, \Pi_2, S, S$ (as Π_1, Π_2 are trivially suitable for any A, due to $S \setminus S_i = \emptyset$ in Df. 6.4). Each atom of S is covered w.r.t. S both by Π_1 and Π_2. As SAT0 is recurrent, by Th. 6.5 each such tree is complete w.r.t. S.

Example 6.8. We continue with program SAT0 and specification S from the previous example, and add a second case of pruning. When the selected atom is of the form $A = q(s_1, s_2)$ with a ground s_1 then only one of clauses (10), (11) is needed – (10) when s_1 is of the form t-t, and (11) otherwise. The other clause can be abandoned without losing the completeness w.r.t. S.[4]

Actually, SAT0 is included in a bigger program, say $P = \text{SAT0} \cup \Pi_0$. We skip the details of Π_0, let us only state that P is recurrent, Π_0 does not contain any clause for p or for q, and that P is complete w.r.t. a specification $S' = S \cup S_0$ where S_0 does not contain any p- or q-atom. (Hence each atom of S_0 is covered by Π_0 w.r.t. S'.)

To formally describe the trees for P resulting from both cases of pruning, consider the following programs and specifications:

$$\begin{array}{ll}
\Pi_1 = \{(7),(8)\}, \;\; \Pi_2 = \{(7),(9)\}, & S_1 = S_2 = S \cap \{\, p(s,u) \mid s,u \in \mathcal{HU} \,\}, \\
\Pi_3 = \{(10)\}, & S_3 = S \cap \{\, q(t{-}t,s) \mid t,s \in \mathcal{HU} \,\}, \\
\Pi_4 = \{(11)\}, & S_4 = S \cap \{\, q(t{-}u,s) \mid t,u,s \in \mathcal{HU}, t \neq u \,\}, \\
\Pi_5 = \{(12)\}, & S_5 = S_=.
\end{array}$$

Let $\mathcal{S}$ be $\Pi_0, \ldots, \Pi_5, S_0, \ldots, S_5$. Each atom from S_i is covered by Π_i w.r.t. S' (for $i = 0, \ldots, 5$). For each q-atom A with its first argument ground, Π_3 or Π_4 is suitable. (Both Π_3, Π_4 are suitable when $ground(A) \cap S = \emptyset$.) For each remaining atom from $\mathcal{TB}$ (at least) one of programs $\Pi_0, \Pi_1, \Pi_2, \Pi_5$ is suitable.

Consider a pruned SLD-tree T for P (employing the two cases of pruning described above). Assume that each q-atom selected in T has its first argument ground. Then T is a csSLD-tree compatible with $\mathcal{S}$. From Th. 6.5 it follows that T is complete w.r.t. S'.

The restriction on the selected q-atoms is implemented by means of Prolog delays. This is done in such a way that, for the intended initial queries, floundering is avoided [Howe and King 2012] (i.e. an atom is selected in each query). So the obtained pruned trees are as T above, and the pruning preserves completeness of the program.

6.2.1. Sufficient condition not applicable directly. Th. 6.5 has an additional consequence. Not only T but also any its subtree is complete w.r.t. S (as the subtree satisfies the premises of Th. 6.5). So the theorem is not (directly) applicable when a csSLD-tree is complete, but pruning removes some answers for a query within the tree (cf. Ex. 6.9 below). In Ex. 6.10 we show how, at least in some such cases, completeness can be proved by employing a family of specifications.

---

[4] The same holds for A of the form $q(s_{11}\text{-}s_{11}, s_2)$, or $q(s_{11}\text{-}s_{12}, s_2)$ with non-unifiable s_{11}, s_{12}. The pruning is implemented using the if-then-else construct in Prolog: `q(V-P,[A|T]) :- V=P -> true ; p(A,T).` (And `p(V-P,[B|T]) :- nonvar(V) -> q(V-P,[B|T]) ; q(B,[V-P|T])` implements the first case of pruning.)

Example 6.9. Consider the logic program P from Ex. 6.2. Predicate c of P is complete w.r.t. a specification

$$S_c = \left\{ c(l_1, l_2) \in \mathcal{HB} \;\middle|\; \begin{array}{l} l_1, l_2 \text{ are lists with} \\ \text{a common element} \end{array} \right\}.$$

The pruning does not impair completeness for ground queries of the form $Q = c(l_1, l_2)$. To apply Th. 6.5 to prove the completeness, each atom of S_c has to be covered by the first clause of P w.r.t. a specification S. Hence $S \supseteq S_c \cup S_m$ is needed where $S_m = \{\, m(t_i, [t_1, \ldots, t_n]) \in \mathcal{HB} \mid n > 0,\ 1 \leq i \leq n \,\}$.

Consider now a pruned tree for such Q, and its subtree T'' rooted at $Q' = m(X, l_1), m(X, l_2)$. The subtree is not complete whenever l_1, l_2 have more than one common element. Thus, for each such case, Th. 6.5 is inapplicable for any superset of S_c.[5]

Example 6.10. The previous example shows a case where completeness of pruned trees w.r.t. a specification cannot be shown by Th. 6.5. However the completeness can be proved by employing a family of specifications. Consider the same program P and specification S_c. For each $c(l_1, l_2) \in S_c$ we construct a specification S' for which P is semi-complete, $c(l_1, l_2) \in S'$, and Th. 6.5 is applicable. Informally, S' may be understood as the subset of $\mathcal{M}_P$ needed to obtain $c(l_1, l_2)$, by means of the clauses involved in a pruned tree for $c(l_1, l_2)$.

Let $l_1 = [t_1, \ldots, t_k]$, $l_2 = [u_1, \ldots, u_n]$. So there exist a term t and indices g, h such that $t = t_g = u_h$, for $i = 1, \ldots, g-1$ no t_i is a member of l_2, and $u_j \neq t$ for $j = 1, \ldots, h-1$. Let $S' = \{\, c(l_1, l_2) \,\} \cup S'_m$ where

$$S'_m = \{\, m(t, [t_i, \ldots, t_k]) \mid 1 \leq i \leq g \,\} \cup \{\, m(t, [u_j, \ldots, u_n]) \mid 1 \leq j \leq h \,\}.$$

Let $\mathcal{S} = \Pi_1, \Pi_2, S_1, S_2$, where Π_i are as in Ex. 6.2, $S_1 = S'$, $S_2 = \{\, m(t, [t_g, \ldots, t_k]), m(t, [u_h, \ldots, u_n]) \,\}$. All the atoms from S_i are covered by Π_i w.r.t. S' (for $i = 1, 2$). Consider the csSLD-tree T' for the (ground) query $Q = c(l_1, l_2)$. From the description in Ex. 6.2 of the c-selection rule, it follows that whenever Π_2 is selected together with an atom A in T' then $A = m(X, [t_g, \ldots, t_k])$ (for some variable X) or $A = m(t, [u_h, \ldots, u_n])$. In both cases $ground(A) \cap S' \subseteq S_2$. Thus T' is compatible with $\mathcal{S}$ and, as it is finite, by Th. 6.5 T' is complete w.r.t. S'.

We showed that, for each $Q \in S_c$, the pruned tree is complete w.r.t. $\{Q\}$. Hence for each ground query of the form $Q = m(s_1, s_2)$ the csSLD-tree for Q is complete w.r.t. S_c.

6.2.2. Another sufficient condition for completeness. By generalizing Th. 5.10 to pruned SLD-trees, we obtain a sufficient condition for completeness, which does not require that the trees are finite. See Appendix B for a proof.

PROPOSITION 6.11. *Let $P \supseteq \bigcup_{i=1}^{n} \Pi_i$ $(n > 0)$ be a program, $S = \bigcup_{i=1}^{n} S_i$ a specification, $|\ |\colon \mathcal{HB} \hookrightarrow \mathbb{N}$ a level mapping, and T a csSLD-tree for $\Pi_1, \ldots, \Pi_n$. If*

(1) T is compatible with $\Pi_1, \ldots, \Pi_n, S_1, \ldots, S_n$,
(2) each atom from S_i is recurrently covered by Π_i w.r.t. S and $|\ |$, for each $i = 1, \ldots, n$,

then T is complete w.r.t. S.

Example 6.12. Consider the program P, specification S and the level mapping $|\ |$ from Ex. 5.11. Consider a c-selection rule R which selects the first atom of the query, and if the

[5] It is not applicable in a more general case of l_1, l_2 having a common element, which is not the last element of l_1. The pruned tree for such Q contains a node with $A = m(Y, [t_i, \ldots, t_k])$ and $\Pi_2 = \{\, m(X, [X|_]).\}$ selected (where $t_i, \ldots, t_k$ are ground, and t_i is a member of l_2). For such atom A, a suitable specification S_2 is a superset of $ground(A) \cap S \supseteq \{\, m(t_j, [t_i, \ldots, t_k]) \mid i \leq j \leq k \,\}$. But such S_2 cannot be covered by Π_2, unless all $t_i, \ldots, t_k$ are equal. So the premises of Th. 6.5 are not satisfied.

Note that this reasoning holds for any representation of the considered pruned trees as csSLD-trees, possibly employing some other $n, \Pi_1, \ldots, \Pi_n$.

atom is of the form $p(t,t)$ it selects $\Pi_1 = \{\, p(X,X) \,\}$, if it is of the form $p(t,u)$ with distinct ground t,u then it selects $\Pi_2 = \{\, p(X,Z) \leftarrow e(X,Y),\ p(Y,Z) \,\}$, otherwise the whole program is selected, $\Pi_3 = P$. Let $S_1 = \{\, p(t,t) \mid t \text{ is a node of } E \,\}$, $S_2 = \{\, p(t,u) \in S \mid t \neq u \,\}$, $S_3 = S$. Note that R is compatible with $\mathcal{S} = \Pi_1, \Pi_2, \Pi_3, S_1, S_2, S_3$ (as in the first case $ground(p(t,t)) \cap S \subseteq S_1$, and in the second case $\{p(t,u)\} \cap S \subseteq S_2$).

From Ex. 5.11 it follows that each atom from S_i is recurrently covered by Π_i w.r.t. S, for $i = 1, 2, 3$. Thus by Prop. 6.11, any csSLD-tree for P via R is complete w.r.t. S. Note that the tree may have infinite branches.

Assume that the initial queries are ground. Then, under the Prolog selection rule, all the selected p-atoms are ground. So the c-selection rule R can be implemented in Prolog by adding the cut to the first clause of P: `p(X,X):-!.` However when the csSLD-tree is infinite, not the whole tree may searched by Prolog due to the depth-first search strategy.

6.3. Correctness of answers of pruned trees

The answers of pruned trees are, obviously, logical consequences of the program. However, it may happen that, due to pruning, not all such consequences are obtained. This feature is sometimes used in Prolog programming; it is called "red cuts" [Sterling and Shapiro 1994]. In this section we show how to characterize more precisely the answers in such cases. We present a sufficient condition for the answers of csSLD-trees to be correct (cf. Df. 2.4), w.r.t. a specification for which the program may be not correct. Informally, the idea is to apply the condition of Th. 4.1 only to those instances of program clauses which were employed in the pruned trees.

THEOREM 6.13. *Let T be a csSLD-tree for a query Q and programs $\Pi_1, \ldots, \Pi_n$, and S be a specification. If for each $i = 1, \ldots n$*

> *$B_1, \ldots, B_m \in S$ implies $H \in S$, for each $(H \leftarrow B_1, \ldots, B_m) \in ground(\Pi_i)$ $(m \geq 0)$*
> *where H is an instance of some atom A selected together with Π_i in T,*

then each answer $Q\theta$ of T is correct w.r.t. S (i.e. $S \models Q\theta$).

PROOF. Assume the premises of the theorem. We have to show that $S \models Q\theta$. Answer $Q\theta$ is obtained out of a successful branch of the tree. The branch represents a derivation with queries $Q_0, \ldots, Q_k$, mgu's $\theta_1, \ldots, \theta_k$, and renamed clauses $C_1, \ldots, C_k$, where Q_k is empty, $Q_0 = Q$, and $Q\theta = Q\theta_1 \cdots \theta_k$. By induction on $k - i$ we obtain $S \models Q_i\theta_{i+1} \cdots \theta_k$ for $i = k, k-1, \ldots, 0$. The base case is obvious. For the inductive step, assume $S \models Q_i\theta_{i+1} \cdots \theta_k$. Let $Q' \in ground(Q_{i-1}\theta_i \cdots \theta_k)$. Let an atom A and a program Π_j be selected for Q_{i-1}. So $C_i \in \Pi_j$. A ground instance Q'' of $Q_i\theta_{i+1} \cdots \theta_k$ can be obtained out of Q' by replacing an instance H of A by a sequence of atoms $B_1, \ldots, B_m$ $(m \geq 0)$, where $(H \leftarrow B_1, \ldots, B_m) \in ground(C_i)$. As each atom of Q'' is in S, so is each atom of Q'. □

Example 6.14. Let us continue Ex. 6.6. The program is obviously not correct w.r.t. S. However, we show that for queries of the form $A = nop(t, Y)$, where t is ground, the answers of pruned trees are correct w.r.t. S.

The condition of Th. 6.13 obviously holds for Π_1, Π_2. Consider Π_3, it is selected together with $A = nop(t, Y)$ only if $t \notin \{adam, eve\}$. Then any ground instance H of both A and the head $nop(X, 2)$ of the clause from Π_3 is in S. Thus the premise of Th. 6.13 is satisfied.

On the other hand, it is not satisfied for e.g. a query $A' = nop(Y, 2)$ (due to its instance $H = nop(adam, 2) \notin S$, $H \in ground(\Pi_3)$ and Π_3 is selected together with A'). Note that the pruned tree for such A' is not correct w.r.t. S.

7. RELATIONS WITH DECLARATIVE DIAGNOSIS

We now discuss the relationship between program diagnosis, and proving correctness and completeness of programs. We first introduce declarative diagnosis of logic programs, and

compare it with the proof methods of Sections 4, 5; details irrelevant for this paper are omitted. Then we show how diagnosis can be performed by attempting a correctness or completeness proof of the program concerned. Finally we discuss a major drawback of declarative diagnosis, and show how it can be avoided by employing approximate specifications.

7.1. Declarative diagnosis

Declarative diagnosis, called sometimes algorithmic, or declarative, debugging, was introduced by Shapiro [1983] (see also [Drabent et al. 1989; Naish 2000] and references therein). It provides methods to locate the reasons of incorrectness or incompleteness of programs. A diagnosis algorithm starts with a *symptom*, obtained from testing the program; the symptom shows incorrectness or incompleteness of the program w.r.t. its specification S. The symptom is an answer Q such that $S \not\models Q$, or a query Q for which the computation terminates but some answers required by S are not produced. (A less general notion of incompleteness symptom is an atom $A \in S$ for which the program finitely fails.) The located error turns out to be a fragment of the program which violates our sufficient condition for correctness or, respectively, semi-completeness.

Incorrectness diagnosis. In declarative diagnosis, the reason for incorrectness is an incorrect instance of a program clause. An incorrect clause instance is one which violates the sufficient condition of Th. 4.1. Obviously, by Th. 4.1 if the program is incorrect then such clause must exist. On the other hand, a violation of the sufficient condition does not imply that the program is incorrect (cf. Ex. A.1 in the Appendix). Incorrectness diagnosis cannot be made more precise; we cannot say which part of the incorrect clause is wrong, as different modifications of various parts of the incorrect clause may lead to a correct one.

Technically, an incorrectness diagnosing algorithm works by constructing a proof tree for an atomic symptom Q, and then searching the tree for an incorrect clause instance. Notice that the actions performed by the algorithm boil down to checking the sufficient condition for correctness of Th. 4.1, but only for some clause instances – those involved in producing the symptom.

Example 7.1 (*Incorrectness diagnosis*). Consider a buggy insertion sort program [Shapiro 1983, Sect. 3.2.3]:

```
isort([],[]).
isort([X|Xs],Ys) :- isort(Xs,Zs), insert(X,Zs,Ys).

insert(X,[],[X]).
insert(X,[Y|Ys],[Y|Zs]) :- Y > X, insert(X,Ys,Zs).
insert(X,[Y|Ys],[X,Y|Ys]) :- X =< Y.
```

(where `>`, `=<` are arithmetic built-ins of Prolog [Apt 1997]). Procedure *isort* should define the sorting relation for lists of integers. Procedure *insert* should describe inserting an element into a sorted list, producing a sorted list as a result. The program computes an incorrect answer $A = isort([2,1,3],[2,3,1])$. (Formally, $S \not\models A$, where S is a specification known by the user; we do not formalize S here.) Thus A is a symptom of incorrectness. The diagnoser constructs a proof tree

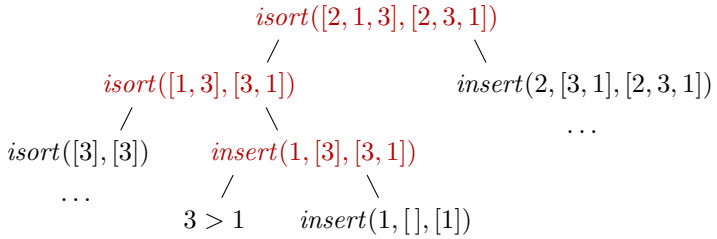

Starting from the root of the proof tree, the diagnosis algorithm asks the user whether the children of the current node are correct w.r.t. the user's specification. When all the answers are YES, the current node together with its children are an incorrect clause instance – the error is found. An answer NO means that a new incorrect node is found, and it becomes the current node.

For the tree above, the questions are first asked about the two children of the root. The child $A' = isort([1,3],[3,1])$ is found incorrect, formally $S \not\models A'$. (The question about the second child is discussed in the next example.) Queries about the children of A' lead to identifying $insert(1,[3],[3,1])$ as incorrect. As both its children are found correct, the algorithm returns $insert(1,[3],[3,1]) \leftarrow 3>1, insert(1,[\,],[1])$ as an incorrect clause instance. (The nodes found incorrect are marked in the diagram.)

Note that more precise locating of the error is impossible, The wrong clause may be corrected for instance by swapping X and Y in $Y > X$, or in the first and last atom of the clause. Thus we cannot determine which fragment of the clause is erroneous.

Incompleteness diagnosis. There exist a few versions of incompleteness diagnosis algorithms ([Naish 1992] provides a comparison). Some of algorithms start with a symptom which is a finitely failed specified atom $A \in S$. Some use, instead, a query Q for which the SLD-tree is finite, and some answers required by the specification S are missing. As the reason for incompleteness, a not covered specified atom is found, say $p(\ldots) \in S$. Alternatively, the result of the diagnosis is an atom $B = p(\ldots)$ which has a not covered specified instance $B\theta \in S$. In both cases procedure p is identified as erroneous.

The existence of a not covered specified atom violates the sufficient condition for semi-completeness of Th. 5.6. Conversely, if the program is not complete for a query Q with a finite SLD-tree then it is not semi-complete and, by Th. 5.6, there must exist a not covered specified atom. Again, violating the sufficient condition for completeness does not imply that the program is not complete. Roughly speaking, the specification may be not sufficiently general, cf. Ex. A.1. Similarly to incorrectness, incompleteness diagnosis cannot be made more precise: A whole procedure is located, and it cannot be determined which fragment of the procedure is wrong, or whether a clause is missing.

We skip an example, and details of incompleteness diagnosis algorithms, see e.g. [Shapiro 1983; Drabent et al. 1989; Naish 1992]. We only point out a relation between completeness proving and some of such algorithms (like that of [Shapiro 1983]). The actions performed by such algorithm boil down to checking the sufficient condition for semi-completeness of Th. 5.6. In contrast to completeness proofs, the condition is checked only for some specified atoms, they are instances of atoms selected in the SLD-tree producing the symptom.

Diagnosis by proof attempts. In a sense, the proof methods presented in this paper supersede the declarative diagnosis methods. An attempt to prove a buggy program to be correct (complete) results in violating the corresponding sufficient condition for some clause (specified atom). Any error located by diagnosis will also be found by a proof attempt; moreover no symptom is needed, and all the errors are found. For instance, in this way the author found an error in an early version of one of the examples from [Drabent 2012a]. However the sufficient condition has to be checked for all the (instances of the) clauses of the program (when proving correctness), or for all specified atoms (in the case of completeness). In contrast, diagnosis takes into account only the clauses (respectively, atoms) related to the computation that produced the symptom.

7.2. A main drawback of declarative diagnosis

Declarative diagnosis requires the user to know an exact specification (a single intended model) of the program. The program has to be (made) correct and complete w.r.t. this specification. This is a serious difficulty in diagnosing actual programs [Drabent et al. 1989, Sect. 26.8], [Naish 2000]. Some diagnoser queries, like "is $append([a],b,[a|b])$ correct", may

be confusing and hardly possible to answer, as the programmer often does not know some details of the intended model, like those related to applying *append* to non lists. In the author's opinion this was a main reason of the lack of acceptance of declarative diagnosis in practice.

Example 7.2 (*Exact specification unfeasible*). The previous example provides a diagnoser query which is hardly possible to answer. Determining correctness of $B = insert(2,[3,1],[2,3,1])$ is confusing, as the user does not know how insertion into unsorted lists should behave. From a formal point of view, the user does not know an exact specification of *insert*.

Actually, both cases of atom B being correct or incorrect are reasonable. They result in locating various errors, and lead to various versions of the final, corrected program. If the specification for *insert* is [6]

$$S^0_{insert} = \left\{ insert(n,l_1,l_2) \in \mathcal{HB} \;\middle|\; \begin{array}{l} l_1, l_2 \text{ are sorted lists of integers,} \\ elms(l_2) = \{n\} \cup elms(l_1) \end{array} \right\},$$

where $elms(l)$ is the multiset of the elements of list l, then $S^0_{insert} \not\models B$ and an erroneous clause instance is eventually found within the subtree rooted in B. Correcting the program would lead to a procedure *insert* which is correct w.r.t. S^0_{insert} but inefficient, as it assures that two of its arguments are sorted lists of integers.

If the specification is

$$S_{insert} = \left\{ insert(n,l_1,l_2) \in \mathcal{HB} \;\middle|\; \begin{array}{l} \text{if } n \in \mathbb{Z} \text{ and } l_1 \text{ is a sorted lists of integers,} \\ \text{then } insert(n,l_1,l_2) \in S^0_{insert} \end{array} \right\}$$

then $S_{insert} \models B$, and the diagnosis continues as in Ex. 7.1. Each of the alternative corrections described there leads to a program with *insert* correct w.r.t. S_{insert}. One of the programs is the "right" one – Program 3.21 of [Sterling and Shapiro 1994] (with $insert(X,[Y|Ys],[Y|Zs]) \leftarrow X > Y, insert(X,Ys,Zs)$ as the corrected clause).

Notice that neither S_{insert} nor S^0_{insert} is an exact specification of procedure *isort* in the "right" program. Providing its exact specification is not so obvious. The "right" procedure is correct w.r.t. S_{insert}, and complete w.r.t. S^0_{insert}. Any such procedure, combined with procedure *isort* from Ex. 7.1, would implement a correct and complete predicate *isort*. The programmer should have the freedom of choosing (and changing) the actual semantics of *insert*, guided by e.g. the efficiency of the program. An exact specification would be counterproductive, at least at earlier stages of program development.

So for this procedure one should take S_{insert} as the specification for correctness and S^0_{insert} as the specification for completeness, and apply them respectively in incorrectness and incompleteness diagnosis.

The discussion in the example suggests a natural way to overcome the difficulty due to declarative diagnosis being based on a single intended model of a program:

Remark 7.3. Declarative diagnosis should employ approximate specifications. In the incorrectness diagnosis the specification for correctness should be used, and in the incompleteness diagnosis – the specification for completeness.

The problem with exact specifications was first dealt with in [Pereira 1986]. In that approach a specification also describes which atoms are *admissible goals*, i.e. are allowed to be selected in SLD-trees. Thus specifications and diagnosis are not declarative, as they refer

[6] An (exact) specification for the remaining predicates is obvious:
$\left\{ sort(l_1,l_2) \in \mathcal{HB} \;\middle|\; \begin{array}{l} l_1 \text{ is a list of integers,} \\ l_2 \text{ is a sorted permutation of } l_1 \end{array} \right\} \cup \left\{ \begin{array}{r} >(i_1,i_2) \in \mathcal{HB}, \\ =<(j_1,j_2) \in \mathcal{HB} \end{array} \;\middle|\; \begin{array}{l} i_1,i_2,j_1,j_2 \text{ are integers,} \\ i_1 > i_2,\ j_1 \leq j_2 \end{array} \right\}.$

to the operational semantics. Naish [2000] introduces a three valued approach, in which a specification classifies each ground atom as *correct*, *erroneous* or *inadmissible*. For such specifications, three-valued declarative debugging algorithms are presented. From our point of view, the set of non-erroneous atoms can be understood as a specification for correctness, and the set of correct atoms as a specification for completeness. However introducing debugging algorithms based on a three-valued logic seems to be an unnecessary complication. In the approach proposed here the standard two-valued diagnosis algorithms remain unchanged.

8. RELATED WORK

Approximate specifications. The difficulties caused by a restriction to exact specifications have been recognized by a few authors. Naish and Søndergaard [2014, and the references therein] propose employing a 3-valued or 4-valued logic. The logical value **i** (inadmissible) is given to the atoms that, in our approach, are in the difference $S_{corr} \setminus S_{compl}$ between the specification for correctness and that for completeness.

Apt [1997] and Deville [1990] present approaches in which it is possible to deal with a suitable subset of $\mathcal{M}_P$ instead of $\mathcal{M}_P$ itself. Apt [1997] employs preconditions and postconditions,[7] and a sufficient condition that, for any atomic query A, if $A \in pre$ then $A\theta \in post$ for any computed answer $A\theta$ (where *pre* and *post* is, respectively, a pre- and postcondition). In this way $\mathcal{M}_P$ is in a sense replaced by *post*, more precisely by the set $\mathcal{M}_P \cap pre$, called $\mathcal{M}_{(pre,post)}$. (See also *Proof methods* below.) Eg. for program APPEND the precondition may be $pre = \{\, app(k,l,m) \in \mathcal{TB} \mid l \text{ or } m \text{ is a list} \,\}$, and the postcondition may describe exactly the list concatenation relation: $post = S^0_{\text{APPEND}}$, in the notation of Ex. 3.1. This gives $\mathcal{M}_{(pre,post)} = S^0_{\text{APPEND}}$. See [Drabent and Miłkowska 2005] for further comparison.

In [Deville 1990] the notions of correctness and completeness of programs involve, for each predicate of the program, the intended relation (on ground terms) and the so-called domain. Let us represent them as Herbrand interpretations, the intended relations as $rel \subseteq \mathcal{HB}$ and the domains as $dom \subseteq \mathcal{HB}$. Then correctness (called there partial correctness) of a definite clause program P is equivalent to $\mathcal{M}_P \cap dom \subseteq rel \cap dom$, and completeness means inclusion in the other direction. For the APPEND example, $rel = S^0_{\text{APPEND}}$ and $dom = \{\, app(k,l,m) \in \mathcal{HB} \mid k,l,m \text{ are lists} \,\}$.

Sometimes it is claimed that programs like APPEND should be modified, so that the resulting program defines "the right" relation [Naish 1996; Apt 1993; Apt 1995]. For the case of APPEND, the new program should define exactly the list concatenation. This can be done by modifying its non-recursive clause so that its variable is restricted to be bound to a list. The clause becomes $app([\,], L, L) \leftarrow list(L)$, and an appropriate definition of *list* is added. The program is believed to be simpler to reason about, but it is impractical due to its inefficiency. Naish [1996, p. 187] proposes that the meaning of the original program should be defined as the standard meaning of the modified program.

The reader is encouraged to compare the treatment of APPEND in Examples 3.1, 4.3, 5.8 to that in the approaches outlined here. To sum up, for reasoning about correctness and completeness of definite programs, the approach based on approximate specifications seems preferable, as it is simpler than the approaches outlined above, refers only to the the basic notions of logic programming, and seems at least as powerful.

Proof methods. Below we review previous work on proving correctness and completeness of logic programs. We omit most of approaches based on operational semantics.

Our definitions of correctness and completeness (Df. 2.1) are basically the same as those in [Sterling and Shapiro 1994]. However no particular method of reasoning on correctness and completeness is given there, and the example informal proofs are made by analyzing possible proof trees. The correctness proving method of [Clark 1979] used here (Th. 4.1)

[7] Technically, pre- and postconditions are sets of atoms closed under substitution.

should be well-known, but is often neglected. For instance, an important monograph [Apt 1997] uses a method,[8] which is more complicated and not declarative. It proves a certain property of LD-derivations, from which the declarative property of program correctness follows. See [Drabent and Miłkowska 2005] for comparison and argumentation that Clark's method is sufficient, and for its generalization to programs with negation.

Surprisingly little has been done on reasoning about completeness. For instance, it is absent from [Apt 1997][9] and [Hogger 1984].[10] Kowalski [1985] discusses completeness, but the example proofs concern only correctness. As a sufficient condition for completeness of a program P he suggests $P \vdash T_S$, where T_S is a specification in a form of a logical theory. (A symmetrical criterion $T_S \vdash P$ is used for correctness, notice its similarity to that of Th. 4.1.) The condition seems impractical as it fails when T_S contains auxiliary predicates, not occurring in P. It also requires that all the models of P (including $\mathcal{HB}$) are models of the specification. But it seems that such specifications often have a substantially restricted class of models, maybe a single Herbrand model, cf. [Deville 1990].

Stärk [1998] presented an elegant method of reasoning about a broad class of properties of programs with negation, executed under LDNF-resolutions. A tool to verify proofs mechanically was provided. The approach involves a rather complicated induction scheme, so it seems impossible to apply the method informally by programmers.

Deville [1990] introduces a systematic approach to constructing programs. Correctness and completeness of a program follows from construction. No direct sufficient criteria, applicable to arbitrary programs, are given. The underlying semantics of Prolog programs is that of LDNF-resolution. Note that the last two approaches are based on operational semantics, and depend on the order of literals in clauses. So they are not declarative.

Deransart and Małuszyński [1993] present criteria for definite program completeness, in a sophisticated framework of relating logic programming and attribute grammars. We present, as Th. 5.10, their sufficient (and complete) criterion for completeness in a simplified setting and with a short direct proof.

The method introduced here (Th. 5.6 with Prop. 5.4) is a simplification of that from [Drabent and Miłkowska 2005]. Due to a restriction to specifications over Herbrand domains, we could deal here with $ground(P)$ and covered ground atoms, instead of a theory ONLY-IF(P) and its truth in an interpretation, as in the former work. Another difference is employing here a new notion of semi-completeness. Also, the former work uses a slightly different notion of completeness (cf. a remark in Sect. 2), and provides a generalization for programs with negation, with a declarative semantics of [Kunen 1987].

Some ideas of the current work appeared in [Drabent 2012a; 2012b] (see also *Applications* in Sect. 9). This includes a weaker version of Th. 6.5 on completeness of pruned SLD-trees. The author is not aware of any other former work on proving completeness in presence of pruning, or on proving correctness of the answers of pruned SLD-trees (w.r.t. specifications for which the program is not correct). Part of this work was presented in a conference paper [Drabent 2015a].

For the related work on declarative diagnosis see Sect. 7. Dealing with the occur-check problem (see [Apt 1997; Deransart and Małuszyński 1993] and the references therein) is outside of the scope of this paper.

---

[8] It is basically the method of [Bossi and Cocco 1989], (stemming from [Drabent and Małuszyński 1988], a paper focused on non-declarative properties of programs). The employed notion of correctness says that, given a precondition *pre* and a postcondition *post*, if the initial atomic query A is in *pre* then in each LD-derivation each selected atom is in *pre* and each corresponding computed answer is in *post*. Note that it depends on the order of atoms in a clause. It implies a declarative property that, in our terminology, the program is correct w.r.t. $(\mathcal{HB} \setminus pre) \cup (post \cap \mathcal{HB})$.

[9] Instead, for a program P and an atomic query A, a characterization of the set of computed instances of A is studied, in a special case of the set being finite and the answers ground [Apt 1997, Sect. 8.4]. This is based on finding the least Herbrand model (of P or of a certain subset of $ground(P)$).

[10] The notion of completeness is defined in [Hogger 1984], but no sufficient condition is discussed.

9. REMARKS

Declarativeness. In this paper we are mainly interested in declarative reasoning about logic programs. Correctness and completeness of logic programs are declarative properties. However non declarative methods for dealing with correctness are often suggested, cf. the previous section. If it were necessary to reason non declaratively about correctness or completeness of programs then logic programming would not deserve to be called a declarative programming paradigm. This paper presents declarative sufficient conditions for correctness and semi-completeness (Th. 4.1, 5.6). Some of the conditions for completeness are also declarative (Prop. 5.4(2), Th. 5.10).

The remaining sufficient conditions for completeness of Prop. 5.4 are not declarative. Condition (1) involves termination, condition (3) depends on the order of atoms in clause bodies (and implies termination). Notice that declarative completeness proofs either imply termination (when the program is shown to be recurrent, Prop. 5.4(2)), or require reasoning similar to that in termination proofs (when showing that covered atoms are recurrently covered, Th. 5.10). Notice also that in most practical cases termination has to be established anyway. So proving completeness in two steps – a declarative proof of semi-completeness and a proof of termination – may be a reasonable compromise between declarative and non-declarative reasoning. Moreover, in many cases termination proofs can be obtained automatically (see [Codish and Taboch 1999; Mesnard and Bagnara 2005; Schneider-Kamp et al. 2009; Nguyen et al. 2011] and references therein). Semi-completeness alone may be a useful property, as it guarantees that whenever the computation terminates, all the answers required by the specification have been computed.

Note that the sufficient conditions of Section 6 for completeness of pruned trees and correctness of their answers also are declarative, except for conditions 3a, 3c of Th.6.5 (which are similar to conditions (1) and (3) of Prop. 5.4, discussed above). Proving condition 3a, i.e. termination, can be partly automated, as mentioned above.

Granularity of proofs. The sufficient conditions presented in this paper impose a certain granularity of proofs. Correctness proofs deal with separate program clauses. Proofs of semi-completeness deal with whole program procedures (to check that an atom $p(\ldots)$ is covered one has to consider all the clauses for p). For completeness, a certain level mapping has to be found, or program termination is to be considered. In both cases the whole program has to be taken into account.

Interpretations as specifications. Specifications which are interpretations are somewhat limited [Drabent and Miłkowska 2005]. Some program properties cannot be expressed. The problem also concerns e.g. the approaches of [Apt 1997], [Hogger 1984], [Kowalski 1985], [Deville 1990], and the declarative diagnosis. The limitation is due to considering completeness w.r.t. a specification which describes a single relation for a predicate symbol. Such specification cannot express that e.g. for a given t there exists a u such that $p(t,u)$ is the program's answer. It has to explicitly state some (one or more) particular u.

In our framework an approximate specification describes a set $\{\, I \mid S_{compl} \subseteq I \subseteq S_{corr} \,\}$ of specifications; its elements are the least Herbrand models of programs compatible with the specification; the set has the least and the greatest element. We have just shown a case where an approximate specification should describe a set with many minimal elements. (A symmetric case – related to correctness – is e.g. describing that at most one of $p(t,u)$, $p(t,s)$ is allowed to be an answer of the program.) A possible solution may be to introduce specifications in a form of logical theories. Such a theory may include axioms like $\forall t \exists u.\, p(t,u)$, and describe a set inexpressible in our framework.

Applications. We agree with the opinion of [Apt 1993] that "unless the verification method is easy and amenable to informal use, it will be ignored". We want to stress the simplicity and naturalness of the sufficient conditions for correctness (Th. 4.1) and semi-completeness

(Th. 5.6, the condition is a part of each discussed sufficient condition for completeness). Informally, the former states that the clauses of a program should produce only correct conclusions, given correct premises. The latter states that each ground atom that should be produced by P can be produced by a clause of P out of atoms which should be produced by P. The author believes that this is a way a competent programmer reasons about (the declarative semantics of) a logic program.

The practical applicability of the methods presented here is illustrated by a larger example in [Drabent 2012a; 2012b]. The example demonstrates a systematic construction of a non-trivial Prolog program (the SAT solver of [Howe and King 2012]). Starting from a formal specification, a definite clause logic program is constructed hand in hand with proofs of its correctness, completeness, and termination under any selection rule. The final Prolog program is obtained by adding control to the logic program (delays and pruning SLD-trees). Adding control preserves correctness and termination. However completeness may be violated by pruning, and by floundering related to delays. By Th. 6.5, the program with pruning remains complete.[11] Reasoning about floundering is outside of the scope of the current paper; non-floundering of this program can be confirmed by a program analysis algorithm [Genaim and King 2008].

The example shows how much of the programming task can be done declaratively, without considering the operational semantics; how "logic" could be separated from "control." It is important that all the considerations and decisions about the program execution and efficiency are independent from those related to the declarative semantics: to the correctness and completeness of the final program. For the role of approximate specifications in this example, see *Comments* in Sect. 3.

Future work. A natural continuation of the work presented here is generalization of the proof methods for programs with negation. There are three semantics to deal with: 3-valued completion semantics of Kunen (Prolog with additional checks for sound negation [Kunen 1987; Doets 1994]), the well-founded semantics (the same with tabulation, as in XSB [Van Gelder et al. 1991; Swift and Warren 2012]), and the answer set semantics [Brewka et al. 2011, and the references therein]. The approach for Kunen semantics could possibly stem from that of [Drabent and Miłkowska 2005]. Also generalization to constraint logic programming and to CHR (constraint handling rules) could be of interest.

For reasoning about pruning, it should be useful to provide methods which directly refer to the pruning constructs (like the cut) used in a program, instead of employing the c-selection rule. An initial attempt was presented in [Drabent 2015a]. An interesting issue is introducing another form of specifications, to overcome the limitation discussed earlier in this section.

An important task is formalization and automatizing of correctness and completeness proofs; a first step is formalization of specifications. On the other hand, further examples of proofs for practical programs are due, as are experiments with teaching programmers to informally use the proof methods in practice.

10. CONCLUSIONS

This paper presents proof methods for proving correctness and completeness of definite clause programs. The method for correctness [Clark 1979] is simple and natural. It should be well-known, but is often neglected. Little work has been done on proving completeness, and this is the main subject of this paper (Sect. 5). A simplification of the approach of [Drabent and Miłkowska 2005] is presented. We introduce a notion of semi-completeness, for which the corresponding sufficient condition deals with program procedures separately,

[11] In the former paper a weaker version of Th. 6.5 has been used, and one case of pruning is discussed informally. A proof covering both cases of pruning is illustrated here in Ex. 6.8.

while for completeness the whole program has to be taken into account. Semi-completeness and termination imply completeness. In practice this means that if a semi-complete program has terminated then it has produced all the answers required by the specification. The presented sufficient condition for semi-completeness corresponds to a natural, intuitive way of reasoning about programs. It is a necessary condition for program completeness (in a sense made precise in Appendix A). We propose a few sufficient conditions for program completeness; all of them involve the above mentioned condition for semi-completeness. Also, sufficient conditions are given for completeness being preserved under SLD-tree pruning (Sect. 6). This is augmented by a sufficient condition for correctness of the answers of pruned SLD-trees (applicable in describing the effects of so called "red cuts" [Sterling and Shapiro 1994]).

Logic programming can't be considered a declarative programming paradigm unless there exist declarative ways of reasoning about program correctness and completeness (i.e. reasoning which abstracts from the operational semantics). Regrettably, non-declarative methods are usually suggested (cf. Sect. 8). The presented methods for proving correctness and semi-completeness are purely declarative (Sect. 9), however some of the sufficient conditions for completeness are not, as they refer, perhaps indirectly, to program termination. Such methods may nevertheless be useful in practical declarative reasoning about programs, as usually termination has to be established anyway.

A larger example of applying the proof methods described here is provided in [Drabent 2012a; 2012b]. That paper presents a construction of a non-trivial program hand in hand with its correctness and completeness proofs. The construction shows how "logic" can be separated from "control"; how the reasoning about correctness and completeness can be separated from that related to the operational semantics, efficiency, etc (cf. *Applications* in Sect. 9).

We point out advantages of approximate specifications (Sect. 3). They are crucial for avoiding unnecessary complications in constructing specifications and proofs. It is often cumbersome and superfluous to exactly describe the relations computed by a program. Approximate specifications are natural in the process of program development: when starting construction of a program, the relations it should compute are often known only approximately. This suggests an extension of the well-known paradigm of program development by transformations that preserve program semantics (see [Pettorossi et al. 2010] for references); one should also consider transformations which preserve correctness and completeness w.r.t. an approximate specification (cf. *Comments* in Sect. 3).

In Sect. 7 we compared the proof methods with declarative diagnosis (algorithmic debugging). Similarity was demonstrated between two ways of locating errors in programs: by means of declarative diagnosis, and by a failure of an attempt to construct a correctness or completeness proof. We have also shown how approximate specifications help to remove a serious drawback of declarative diagnosis.

We argue that the proof methods presented here are simple, and reflect a natural way of declarative thinking about programs (Sect. 9). We believe that they can be actually used – at least at an informal level – in practical programming; this is supported by examples.

APPENDIX

A. COMPLETENESS OF PROOF METHODS

An important feature of a proof method is its completeness. Here we show in which sense the methods of proving correctness (Th. 4.1), semi-completeness (Th. 5.6), and completeness (Th. 5.10) are complete.

Obviously, we cannot require that if a program P is correct (respectively, complete) w.r.t. a specification S then a proof method applied to P and S shows the correctness (completeness). The specification may provide insufficient information.

Example A.1 (*Specifications too weak (strong) for a correctness (completeness) proof*). Program $P = \{\, p(X) \leftarrow q(X).\ q(a).\,\}$ is correct w.r.t. $S = \{\, p(a), q(a), q(b)\,\}$. However the sufficient condition from Th. 4.1 is violated: clause instance $p(b) \leftarrow q(b)$ has its body atom in S and the head not in S. We need a stronger specification $S' = \{\, p(a), q(a)\,\} \subseteq S$, for which P is correct and the sufficient condition holds.

In Ex. 6.3, program Π_1 is complete w.r.t. $S_0 = \{\, q(s^j(0)) \mid j \geq 0\,\}$, but the sufficient condition for semi-completeness of Th. 5.6 does not hold: some atoms of S_0 (actually all) are not covered by P w.r.t. S_0. To prove the completeness, we need a more general specification $S = S_0 \cup \{\, p(b, s^j(0)) \mid j \geq 0\,\}$. Now each atom S is covered by Π_1 w.r.t. S, and the sufficient condition holds. (Similar reasoning applies to the completeness of P and Π_2 w.r.t. S_0.)

Theorems A.2, A.4 below are from [Deransart and Małuszyński 1993].

THEOREM A.2 (COMPLETENESS OF THE METHOD FOR CORRECTNESS). *If a program P is correct w.r.t. a specification S then there exists a stronger specification $S' \subseteq S$ such that the sufficient condition from Th. 4.1 holds for P and S'.*

PROOF. Take $S' = \mathcal{M}_P$. □

THEOREM A.3 (COMPLETENESS OF THE METHOD FOR SEMI-COMPLETENESS). *If a program P is semi-complete w.r.t. a specification S then there exists a more general specification $S' \supseteq S$ such that the sufficient condition from Th. 5.6 holds for P and S'.*

PROOF. For P and each atom $A \in S$ there exists a successful or infinite SLD-derivation. We consider ground instances of such derivations. Let D be a set of derivations for $ground(P)$, containing a successful or infinite derivation for each $A \in S$. Let S' be the set of atoms selected in the derivations of D. Each atom of S' is covered by P. □

Proving completeness of programs based on Prop. 5.4 and Th. 5.6 is not complete. It fails e.g. when a program contains a void clause $A \leftarrow A$. However the method based on Th. 5.10 is complete.

THEOREM A.4 (COMPLETENESS OF THE METHOD FOR COMPLETENESS, TH. 5.10). *If a program P is complete w.r.t. a specification S then there exists a more general specification $S' \supseteq S$, and a level mapping $|\ |$, such that the sufficient condition from Th. 5.10 holds for P, S' and $|\ |$.*

PROOF. For each $A \in S \subseteq \mathcal{M}_P$ there exists a proof tree for $ground(P)$ with the root A. Let $\mathcal{T}$ be the set of such trees and their subtrees. Let S' be the set of their nodes. For a $B \in S'$, let $|B|$ be the minimal height of a tree from $\mathcal{T}$ with the root B. Consider such a minimal height tree. Let $B_1, \ldots, B_n$ be the children of the root B in the tree. So $|B| > |B_i|$ for each i. As $B_1, \ldots, B_n \in S'$ and $B \leftarrow B_1, \ldots, B_n \in ground(P)$, atom B is recurrently covered by P w.r.t. S' and $|\ |$. □

From the last theorem it follows that if P is complete w.r.t. S then some $S' \supseteq S$ is covered by P. Thus the sufficient condition of Th. 5.6 (applied to some $S' \supseteq S$) is a necessary condition for program completeness (w.r.t. S).

B. PROOFS

This appendix contains the remaining proofs (for sections 2, 5 and 6). Let us note that the results of sections 2–6 hold also for infinite programs, unless stated otherwise.

B.1. The least Herbrand model and the logical consequences of a program

This section contains a proof of Lemma 2.2.

LEMMA 2.2. Let P be a program, and Q a query such that

(1) Q contains exactly $k \geq 0$ (distinct) variables, and the underlying language has (at least) k constants not occurring in P, Q, or a non-constant function symbol not occurring in P, Q.

Then $\mathcal{M}_P \models Q$ iff $P \models Q$.

PROOF. The "if" case is obvious. Assume $\mathcal{M}_P \models Q$. Let $V_1, \ldots, V_k$ be the variables occurring in the query Q. Consider a substitution $\rho = \{V_1/t_1, \ldots, V_k/t_k\}$ where $t_1, \ldots, t_k$ are distinct ground terms whose main function symbols $f_1, \ldots, f_k$ do not occur in $P \cup \{Q\}$. Note that none of $f_1, \ldots, f_k$ occurs in any computed answer for P and Q. As $M_P \models Q\rho$ and $Q\rho$ is ground, $P \models Q\rho$ by Th. 4.30 of [Apt 1997].

By completeness of SLD-resolution, $Q\rho$ is an instance of some computed answer $Q\varphi$ for P and Q: $Q\rho = Q\varphi\sigma$. So (for $i = 1, \ldots, k$) $t_i = V_i\varphi\sigma$, $t_i = f_i(\ldots)$ and f_i does not occur in $V_i\varphi$. Thus each $V_i\varphi$ is a variable. As $t_1, \ldots, t_k$ are distinct, $V_1\varphi, \ldots, V_k\varphi$ are distinct. Thus $Q\varphi$ is a variant of Q and, by soundness of SLD-resolution, $P \models Q$. □

B.2. Sufficient conditions for completeness

Here we prove Th. 5.6, 6.5, and Prop. 6.11. We begin with a slightly more general version of Th. 6.5, from which Th. 6.5 follows immediately. The only difference is condition 3a.

LEMMA B.1. *Let $P \supseteq \bigcup_{i=1}^{n} \Pi_i$ (where $n > 0$) be a program, $S = \bigcup_{i=1}^{n} S_i$ a specification, and T a csSLD-tree for $\Pi_1, \ldots, \Pi_n$. If*

(1) for each $i = 1, \ldots, n$, all the atoms from S_i are covered by Π_i w.r.t. S, and
(2) T is compatible with $\Pi_1, \ldots, \Pi_n, S_1, \ldots, S_n$,
(3) (a) all the branches of T are finite, or
(b) P is recurrent, or
(c) P is acceptable and T is built under the Prolog selection rule

then T is complete w.r.t. S.

PROOF. Let Q' be a node of T which has a ground instance $Q'\sigma$ such that $S \models Q'\sigma$. Let an atom A of Q' and a program Π_i be selected in Q'. Now $A\sigma \in S$, $A\sigma \notin S \setminus S_i$, thus $A\sigma \in S_i$, and $A\sigma$ is covered by Π_i w.r.t. S. Let $A\sigma \leftarrow B_1, \ldots, B_m$ be a ground instance of a clause from Π_i, with $B_1, \ldots, B_m \in S$. Let Q'' be $Q'\sigma$ with $A\sigma$ replaced by $B_1, \ldots, B_m$. Then, by the lifting theorem [Doets 1994, Th. 5.37], Q'' is an instance of a child Q^+ of Q' in T. Obviously, $S \models Q''$.

Assume the root of T is Q. Let $S \models Q\theta$ for a ground instance $Q\theta$ of Q. By induction, using the previous paragraph as the inductive step, there exists a branch Δ in T which (i) is infinite or its last node is the empty query, (ii) each its nonempty node has a ground instance consisting of atoms from S, (iii) the sequence of these ground instances is a derivation Γ for $Q\theta$ and $ground(P)$, and (iv) Δ is a lift [Doets 1994] of Γ.

We show that Δ is finite. This is immediately implied by condition 3a of the Lemma. Condition 3b or 3c implies that $ground(P)$ is recurrent or acceptable. Hence Γ is finite (by [Apt 1997, Corollaries 6.10, 6.25]); thus Δ is finite too. By (i), Δ is a successful derivation (for Q and P), thus so is Γ, by (iv). In the notation of [Doets 1994], the respective answers of Γ, Δ are $Q\theta = res(\Gamma)$ and $res(\Delta)$ (the resultants of the derivations). By the lifting theorem, $Q\theta$ is an instance of $res(\Delta)$, which is an answer of T. □

Now, let us take a program P and a specification S, and let $n = 1$, $\Pi_1 = P$, $S_1 = S$. P is suitable for any atom, as $S \setminus S_1 = \emptyset$. Thus each SLD-tree T for P is a csSLD-tree compatible with P, S, and from Lemma B.1 it follows:

PROPOSITION B.2. *If all the atoms from a specification S are covered w.r.t. S by a program P, and there exists an SLD-tree for P and a query Q with no infinite branches then P is complete for Q.*

Th. 5.6 follows immediately from Prop. B.2. Now we show an example where Th. 5.6 is not applicable, but Prop. B.2 is.

Example B.3. Program $P = \{\, p(s(X)) \leftarrow p(X).\ \ p(0).\ \ q(X) \leftarrow p(s(Y)).\}$ is complete w.r.t. $S = \{p(s^i(0)) \mid i \geq 0\} \cup \{q(0)\}$, but its completeness does not follow from semi-completeness. The program loops for any instance of $Q = q(X)$ (the SLD-tree has an infinite branch). However all the branches of the SLD-tree for $q(X)$ and $ground(P)$ are finite (although the tree has infinite height). Hence by Prop. B.2, $ground(P)$ is complete for Q. Thus so is P, as both programs have the same least Herbrand model.

PROPOSITION 6.11. Let $P \supseteq \bigcup_{i=1}^{n} \Pi_i$ $(n > 0)$ be a program, $S = \bigcup_{i=1}^{n} S_i$ a specification, $|\ |\colon \mathcal{HB} \hookrightarrow \mathbb{N}$ a level mapping, and T a csSLD-tree for $\Pi_1, \ldots, \Pi_n$. If

(1) T is compatible with $\Pi_1, \ldots, \Pi_n, S_1, \ldots, S_n$,
(2) each atom from S_i is recurrently covered by Π_i w.r.t. S and $|\ |$, for each $i = 1, \ldots, n$,

then T is complete w.r.t. S.

PROOF. For a ground query $Q = A_1, \ldots, A_n$, let us define the level of Q as the multiset $|Q| = bag(|A_1|, \ldots, |A_n|)$, whenever $|A_1|, \ldots, |A_n|$ are defined. Let $\prec_m$ be the multiset ordering [Apt 1997].

We follow the proof of Lemma B.1. The first part of the proof is basically the same, with an additional restriction made possible by the atom $A\sigma \in S$ being recurrently covered by Π_i w.r.t. S. Namely for the ground clause $A\sigma \leftarrow B_1, \ldots, B_m$ we additionally require that $|A\sigma|, |B_1|, \ldots |B_n|$ are defined, and and $|A\sigma| > |B_i|$ for all $i = 1, \ldots, n$. Hence $|Q''| \prec_m |Q'\sigma|$.

As in the proof of Lemma B.1, there is a branch Δ in T and a derivation Γ for $Q\theta$ and $ground(P)$, satisfying conditions (i),...,(iv). For each query Q'' of Γ, $S \models Q''$, and $|Q''|$ is defined. As the levels of the queries in Γ are a $\prec_m$-decreasing sequence, and $\prec_m$ is well founded, Γ is finite. Hence Γ and Δ are successful, and $Q\theta$ is an instance of an answer of T. □

REFERENCES

K. R. Apt. 1993. Declarative Programming in Prolog. In *Logic Programming, Proc. of the 1993 International Symposium*, D. Miller (Ed.). The MIT Press, 12–35.

K. R. Apt. 1995. Program Verification and Prolog. In *Specification and Validation Methods for Programming Languages and Systems*, E. Börger (Ed.). Oxford University Press, 55–95.

K. R. Apt. 1997. *From Logic Programming to Prolog.* Prentice-Hall.

K. R. Apt and D. Pedreschi. 1993. Reasoning about Termination of Pure Prolog Programs. *Information and Computation* 106, 1 (1993), 109–157.

M. Bezem. 1993. Strong Termination of Logic Programs. *J. Log. Program.* 15, 1&2 (1993), 79–97.

A. Bossi and N. Cocco. 1989. Verifying Correctness of Logic Programs. In *TAPSOFT, Vol.2 (Lecture Notes in Computer Science)*, J. Díaz and F. Orejas (Eds.), Vol. 352. Springer, 96–110.

G. Brewka, T. Eiter, and M. Truszczyński. 2011. Answer set programming at a glance. *Commun. ACM* 54, 12 (2011), 92–103. DOI:http://dx.doi.org/10.1145/2043174.2043195

K. L. Clark. 1979. *Predicate logic as computational formalism.* Technical Report 79/59. Imperial College, London.

W. F. Clocksin and C. S. Mellish. 1981. *Programming in Prolog.* Springer.

M. Codish and C. Taboch. 1999. A Semantic Basis for the Termination Analysis of Logic Programs. *J. Log. Program.* 41, 1 (1999), 103–123. DOI:http://dx.doi.org/10.1016/S0743-1066(99)00006-0

P. Deransart and J. Małuszyński. 1993. *A grammatical view of logic programming.* The MIT Press. I–XXXI, 1–454 pages.

Y. Deville. 1990. *Logic Programming: Systematic Program Development*. Addison-Wesley.

Y. Deville and K. K. Lau. 1994. Logic Program Synthesis. *J. Log. Program.* 19/20 (1994), 321–350.

K. Doets. 1994. *From Logic to Logic Programming*. The MIT Press, Cambridge, MA.

W. Drabent. 2012a. Logic + control: An example. In *Technical Communications of the 28th International Conference on Logic Programming (ICLP'12) (Leibniz International Proceedings in Informatics (LIPIcs))*, Agostino Dovier and Vítor Santos Costa (Eds.), Vol. 17. 301–311. http://drops.dagstuhl.de/opus/volltexte/2012/3631.

W. Drabent. 2012b. Logic + control: An example of program construction. *CoRR* abs/1110.4978 (2012). http://arxiv.org/abs/1110.4978.

W. Drabent. 2015a. On completeness of logic programs. In *Logic Based Program Synthesis and Transformation, LOPSTR 2014. Revised Selected Papers (Lecture Notes in Computer Science)*, Vol. 8981. Springer. Extended version in *CoRR* abs/1411.3015 (2014). http://arxiv.org/abs/1411.3015.

W. Drabent. 2015b. On definite program answers and least Herbrand models. *CoRR* abs/1503.03324 (2015). http://arxiv.org/abs/1503.03324.

W. Drabent and J. Małuszyński. 1988. Inductive Assertion Method for Logic Programs. *Theoretical Computer Science* 59 (1988), 133–155.

W. Drabent and M. Miłkowska. 2005. Proving correctness and completeness of normal programs – a declarative approach. *Theory and Practice of Logic Programming* 5, 6 (2005), 669–711.

W. Drabent, S. Nadjm-Tehrani, and J. Małuszyński. 1989. Algorithmic Debugging with Assertions. In *Meta-Programming in Logic Programming*, H. Abramson and M. H. Rogers (Eds.). The MIT Press, 501–522.

S. Genaim and A. King. 2008. Inferring Non-Suspension Conditions for Logic Programs with Dynamic Scheduling. *ACM Trans. Comput. Log.* 9, 3 (2008). http://www.cs.kent.ac.uk/pubs/2008/2456

C. J. Hogger. 1984. *Introduction to Logic Programming*. Academic Press, London.

J. M. Howe and A. King. 2012. A Pearl on SAT and SMT Solving in Prolog. *Theor. Comput. Sci.* 435 (2012), 43–55.

R. A. Kowalski. 1985. The Relation between Logic Programming and Logic Specification. In *Mathematical Logic and Programming Languages*, C. Hoare and J. Shepherdson (Eds.). Prentice-Hall, 11–27. Also in *Phil. Trans. R. Soc. Lond. A*, Vol 312, 1984, 345–361.

K. Kunen. 1987. Negation in Logic Programming. *J. Log. Program.* 4, 4 (1987), 289–308.

F. Mesnard and R. Bagnara. 2005. cTI: A constraint-based termination inference tool for ISO-Prolog. *TPLP* 5, 1-2 (2005), 243–257. DOI:http://dx.doi.org/10.1017/S1471068404002017

L. Naish. 1992. Declarative Diagnosis of Missing Answers. *New Generation Comput.* 10, 3 (1992), 255–286.

L. Naish. 1996. A Declarative View of Modes. In *Proceedings of JICSLP'96*. The MIT Press, 185–199.

L. Naish. 2000. A Three-Valued Declarative Debugging Scheme. In *23rd Australasian Computer Science Conference (ACSC 2000)*. IEEE Computer Society, 166–173.

L. Naish. 2006. A three-valued semantics for logic programmers. *TPLP* 6, 5 (2006), 509–538.

L. Naish and H. Søndergaard. 2014. Truth versus information in logic programming. *TPLP* 14, 6 (2014), 803–840. DOI:http://dx.doi.org/10.1017/S1471068413000069

M. T. Nguyen, D. De Schreye, J. Giesl, and P. Schneider-Kamp. 2011. Polytool: Polynomial interpretations as a basis for termination analysis of logic programs. *TPLP* 11, 1 (2011), 33–63. DOI:http://dx.doi.org/10.1017/S1471068410000025

A. Paschke. 2011. Rules and Logic Programming for the Web. In *Reasoning Web. Semantic Technologies for the Web of Data - 7th International Summer School 2011, Galway, Ireland, August 23-27, 2011, Tutorial Lectures (Lecture Notes in Computer Science)*, Axel Polleres *et al.* (Ed.), Vol. 6848. Springer, 326–381. DOI:http://dx.doi.org/10.1007/978-3-642-23032-5_6

L. M. Pereira. 1986. Rational Debugging in Logic Programming. In *ICLP (Lecture Notes in Computer Science)*, E. Y. Shapiro (Ed.), Vol. 225. Springer, 203–210.

A. Pettorossi, M. Proietti, and V. Senni. 2010. The Transformational Approach to Program Development. In *A 25-Year Perspective on Logic Programming: Achievements of the Italian Association for Logic Programming, GULP (Lecture Notes in Computer Science)*, A. Dovier and E. Pontelli (Eds.), Vol. 6125. Springer, 112–135. DOI:http://dx.doi.org/10.1007/978-3-642-14309-0_6

F. Pfenning (Ed.). 1992. *Types in Logic Programming*. The MIT Press.

P. Schneider-Kamp, J. Giesl, A. Serebrenik, and R. Thiemann. 2009. Automated termination proofs for logic programs by term rewriting. *ACM Trans. Comput. Log.* 11, 1 (2009). DOI:http://dx.doi.org/10.1145/1614431.1614433

E. Shapiro. 1983. *Algorithmic Program Debugging*. The MIT Press.

R. F. Stärk. 1998. The Theoretical Foundations of LPTP (A Logic Program Theorem Prover). *J. Log. Program.* 36, 3 (1998), 241–269.

L. Sterling and E. Shapiro. 1994. *The Art of Prolog* (2 ed.). The MIT Press.

T. Swift and D. S. Warren. 2012. XSB: Extending Prolog with Tabled Logic Programming. *TPLP* 12, 1-2 (2012), 157–187. DOI:http://dx.doi.org/10.1017/S1471068411000500

A. Van Gelder, K. A. Ross, and J. S. Schlipf. 1991. The Well-Founded Semantics for General Logic Programs. *J. ACM* 38, 3 (1991), 620–650. DOI:http://dx.doi.org/10.1145/116825.116838